# Simulate the Topology


- Simulate the Topology Overview, page 1
- Cisco Modeling Labs Active Canvas, page 2
- Launch a Simulation, page 5
- Connect to a Simulation Node Console, page 22
- Start a Single Node, page 28
- Stop a Simulation, page 33
- Stop a Single Node, page 38
- Modify a Node Configuration in the Simulation, page 40
- Extract and Save Modified Configurations, page 46
- Linux Server Snapshot Support, page 49
- Latency, Jitter and Packet Loss Control Options, page 52


# Simulate the Topology Overview

The simulation phase is when you run the simulation of your topology design. The **Simulation** perspective provides a set of views that support the simulation phase. By comparison, the design and build activities occur in the **Design** perspective, which provides a set of views that support the design activity. Some views in the **Simulation** perspective can also be viewed in the **Design** perspective.

### Simulation Perspectives and Views

The main areas of focus within the **Simulation** perspective are the **Simulations** view and the **Console** view. The following figure highlights the **Simulations** and **Console** views for a running simulation.

From the **View** menu, you can open additional views in the perspective and arrange the open views by dragging them within the perspective workspace. See the online help and the section Customize Perspectives for more information.

**Note** To reset your current perspective to its original configuration when the workbench was first opened, right-click the perspective button and select **Reset**.

### Node Limits

In Cisco Modeling Labs, the number of nodes you can run is limited by the resources available on the server and the number of licenses available. Currently, the design limit is 1,000 nodes.

# Cisco Modeling Labs Active Canvas

The Cisco Modeling Labs client provides users with an active canvas. When a simulation is started and the user switches to the Simulation perspective, a new window opens displaying the network diagram. As the virtual machines start up, the network diagram updates showing the current state of the simulation. Nodes change color depending on their current operational state.

For example, the following figure shows nodes in green which indicates the **Active** state. Nodes in blue indicate the **Build** state. Grey nodes indicate the **Absent** state where a node is yet to be started or has been stopped.

***Figure 1: Node States in a Simulation***

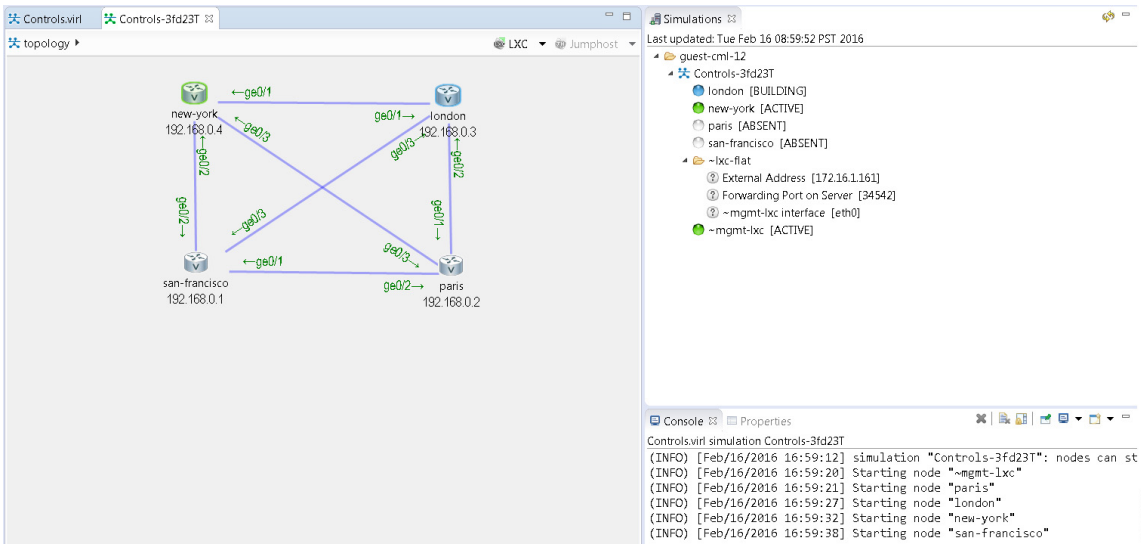

Once a node is in the active state, you can right-click on the node to perform operations such as opening an SSH or Telnet connection, extract the configuration of the specific node and stop/start the node.

***Figure 2: Node Operations***

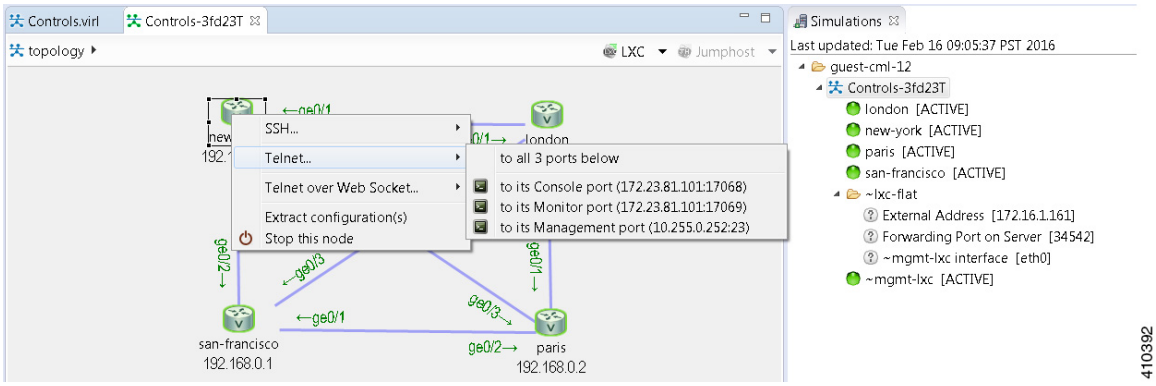


Right-clicking on the background, without selecting a node, enables you to perform simulation-wide operations such as configuration extraction, launch the live visualization view, stop the simulation as well as resetting all link latency, jitter and packet-loss parameters that may be in operation.

***Figure 3: Simulations Operations***

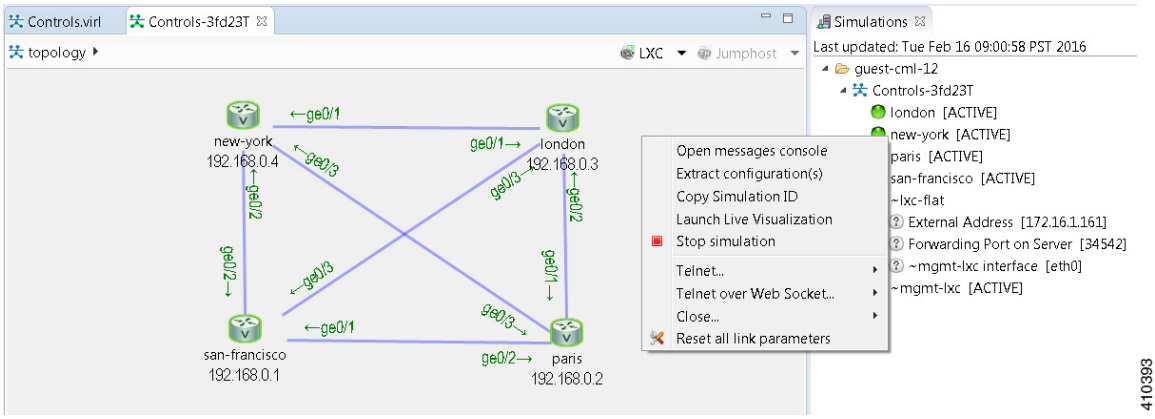

If the **Simulations** view is closed, it can be reopened by selecting the simulation from the simulations panel, right-clicking and selecting the **View Simulation** option.

***Figure 4: View Simulation Option***

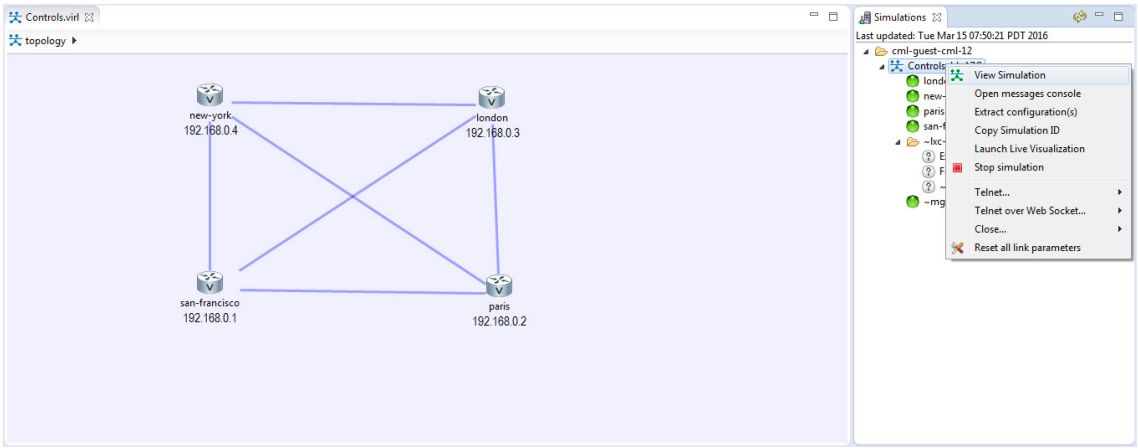


Link latency, jitter and packet-loss parameters can be set by selecting a link, right-clicking and using the **Modify Link Parameters** option.

***Figure 5: Link Parameters***

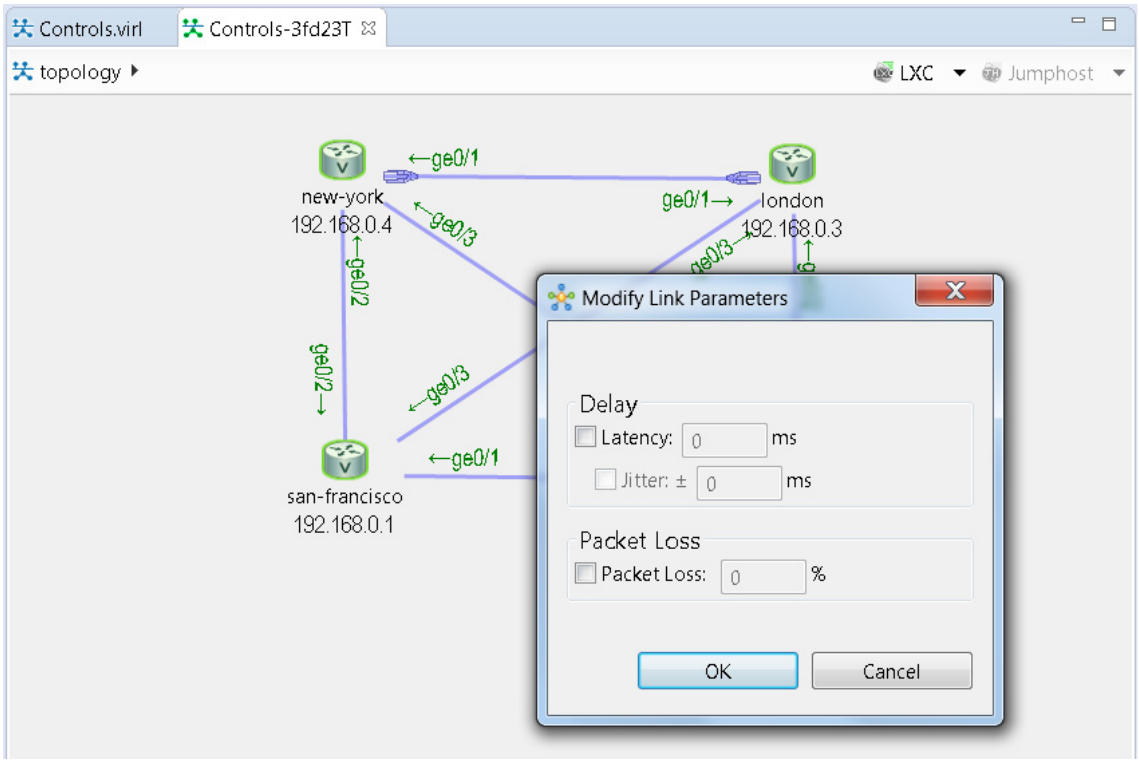

Packet capture operations can be performed by selecting a link, selecting the interface (at one end of the link) and right-clicking to reveal the packet-capture control menu, as shown.

***Figure 6: Packet Capture Operations***

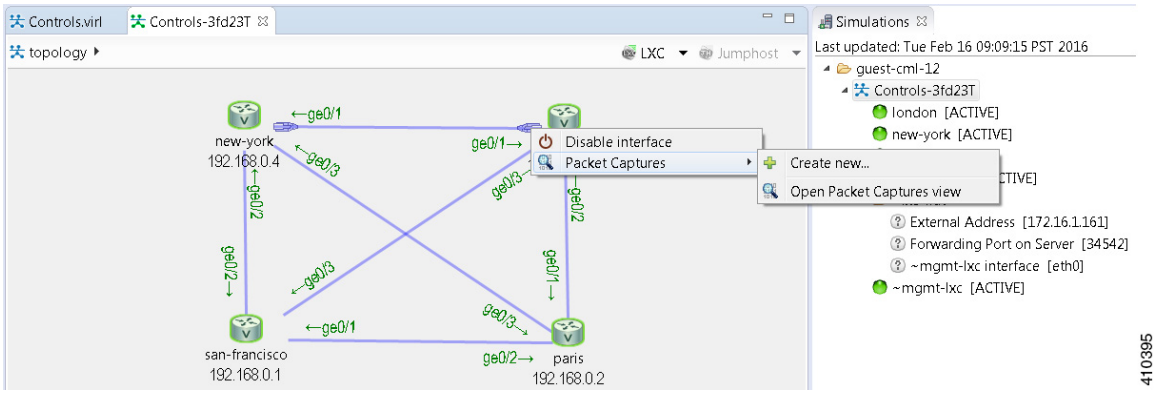


Once a packet capture has been configured, an icon will indicate that a packet-capture is present on the interface, with the Packet capture view listing the .pcap file that is available for analysis.

Additional diagram labels are now available including interface name, serial port number assignment and so on. These can be accessed from the **Show Topology Labels** icon on the Cisco Modeling Labs toolbar.

***Figure 7: Show Topology Label Options***

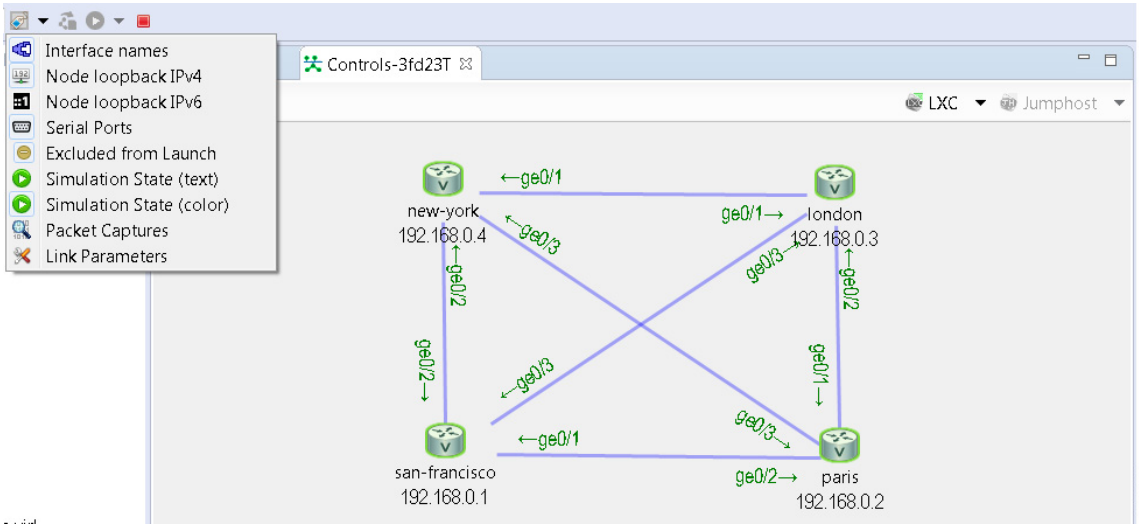


# Launch a Simulation

To launch a simulation, complete the following steps.

## Before You Begin

- Complete the topology design.

- Complete the task of building the nodes and interfaces.
- (Optional) Generate the configuration using AutoNetkit.

**Caution** When you manually make changes to a node configuration and bypass AutoNetkit autogeneration, those changes do not appear in the topology view of the **Design** perspective or **Simulation** perspective. For example, if you use the **hostname** command to change the host name from `iosv-1` to `Router-1` in the configuration, the node name in the topology view and in other related views remains as `iosv-1`.

- Open the desired topology.

**Note** The topology should be open and visible in the **Topology Editor**. If you have multiple topologies open in the **Topology Editor**, simulation will launch for the currently active view.

**Step 1** From the toolbar, click the **Launch Simulation** button.
The simulation launches and provides a unique identifier, which means that multiple instances of the topology can be launched and each will have a unique name.

**Step 2** In the **Launching Simulation** dialog box, select any of the following actions:

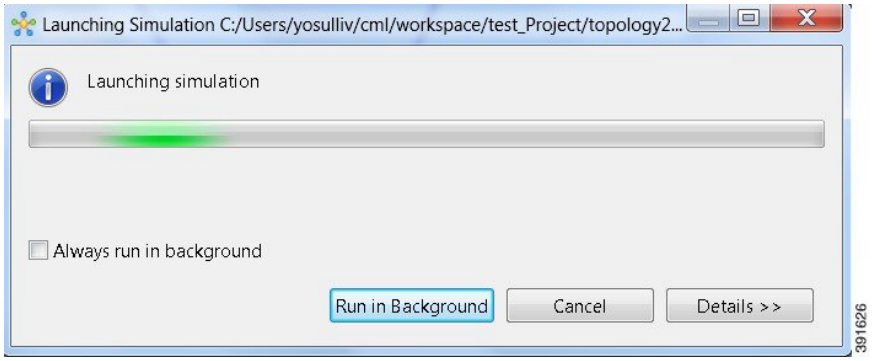


a) Check the **Always run in the background** check box. All future node start requests, stop requests, and simulation launch requests run in the background and do not display dialog boxes.
**Note** To control the background setting, choose **File** > **Preferences** > **General**.

b) Click **Run in Background**. The dialog box closes when the node simulation starts.
c) Click **Cancel** to return to the **Design** view.
d) Take no action, and the node simulation launches momentarily.
**Tip** When you click **Run in Background**, the status bar displays a progress icon. Click the icon to display a compact view of the activity progress. If the background activity encounters an error, the icon shows a red **X**. Click the error icon to display the error dialog.

**Step 3** In the **Simulation launched** dialog box, click **OK**.

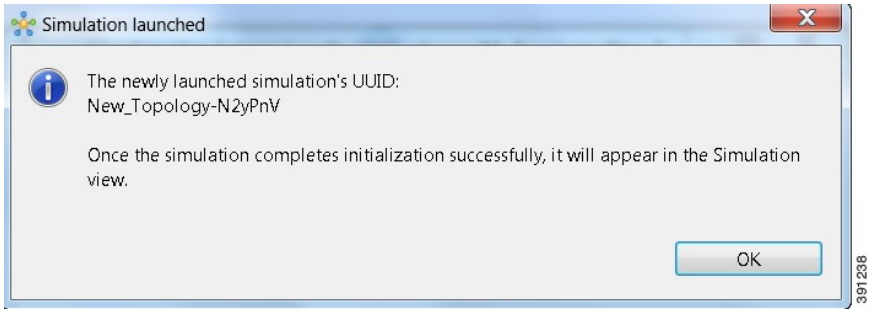


**Step 4** When prompted to switch to the **Simulation** perspective, click **Yes**.

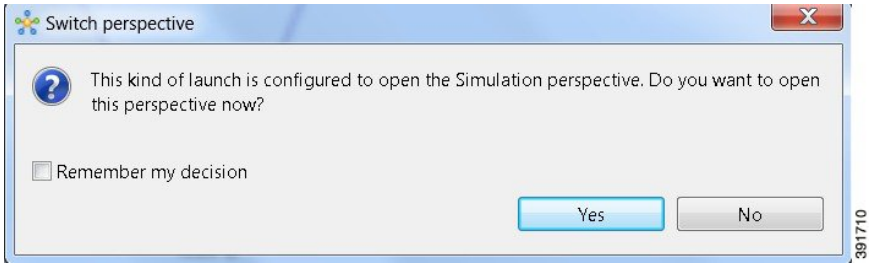


**Note** We recommend that you switch to the **Simulation** perspective to view the running simulation and to connect to node consoles.

The Cisco Modeling Labs client tracks the state of the simulations. All launched simulations appear in the **Simulations** view. Status messages are displayed in the **Console** view. After the Cisco Modeling Labs server has started the simulation launch, a confirmation dialog box appears with an identifier assigned to the simulation by the Cisco Modeling Labs server.

When the nodes are running on the Cisco Modeling Labs server, they are displayed in the **Simulations** view with the status **[ACTIVE]**, as shown in the following figure.

***Figure 8: Simulation Launched in Simulations View***

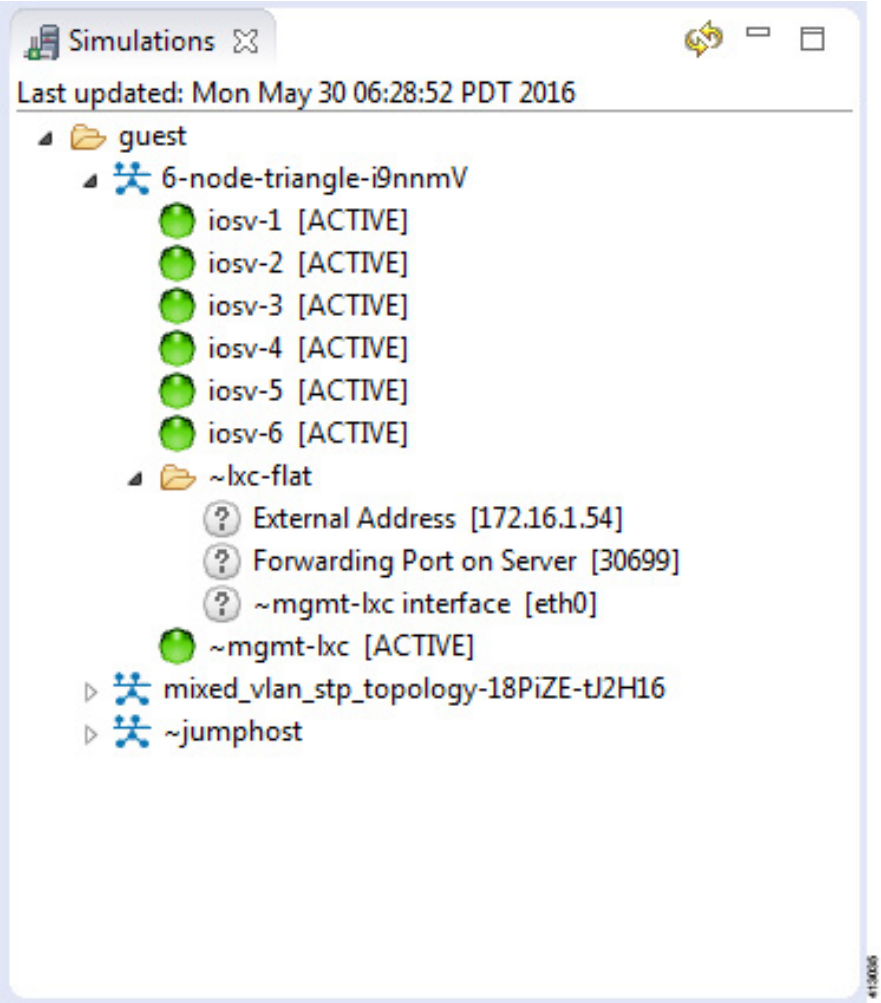


# Jumphost Virtual Machine (VM)

The jumphost VM is the default method for accessing the management network of a running simulation. The jumphost node runs in a separate simulation named **~jumphost** and has two interfaces, eth0 in the project/user management network and eth1 in the FLAT network.

The purpose of the jumphost is to provide an access point into a simulated network that remains fixed, in that there is a single external IP address or port that the user can access. A user can access the jumphost and then access all the nodes inside the simulation.

Cisco Modeling Labs provides two implementations of jumphost:

- A VM: Based on the server VM image type.
- A Linux container: A lighter weight form of a jumphost. See Linux Container (LXC), on page 9 for more information.

The VM implementation is costly in terms of the memory and CPU used when a jumphost is instantiated. However, since it is a full-blown server VM, there is value to it, in that you can install and run any application on it.

To select a jumphost VM, in the Cisco Modeling Labs client, choose **Properties** > **Topology** > **Management Network** > **Private project network** or **Shared flat network**.

***Figure 9: Setting the Jumphost Option***

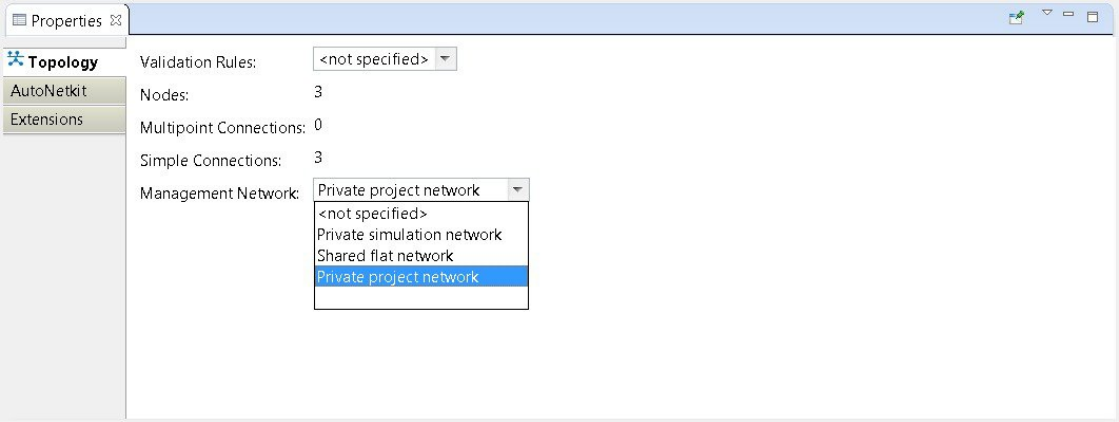


The jumphost VM is per user; a user can create multiple simulations, but only one corresponding jumphost is created. Therefore, depending on the type of simulation you are running, you can choose between the two implementations.

# Linux Container (LXC)

An LXC provides a means of accessing a topology rather than having to create a full Linux server VM. All nodes in the network are connected to a hidden OOB management network that uses the first interface on each of the nodes.

In the Cisco Modeling Labs client, it is enabled by selecting the management network type **Private simulation network** for your topology. The LXC is automatically connected into this hidden OOB management network

to which all VMs in your simulation are connected. This enables you to connect into each VM via its management Ethernet port, removing the need to use the console port connection method.

***Figure 10: Setting the LXC Option***

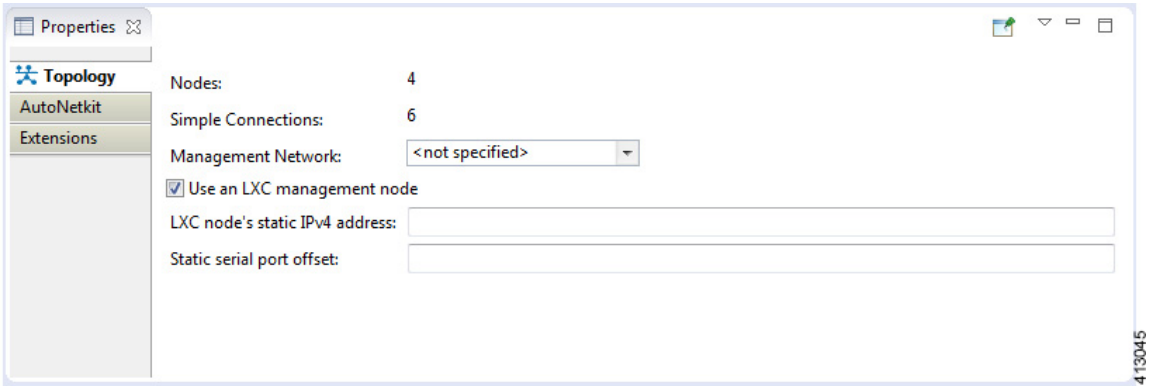


The LXC is operational when your simulation is active and terminates when your simulation stops. LXC uses a significantly smaller memory footprint than the Linux server VM. The LXC automatically gets an IP address on the FLAT network.

***Figure 11: LXC on the FLAT Network***

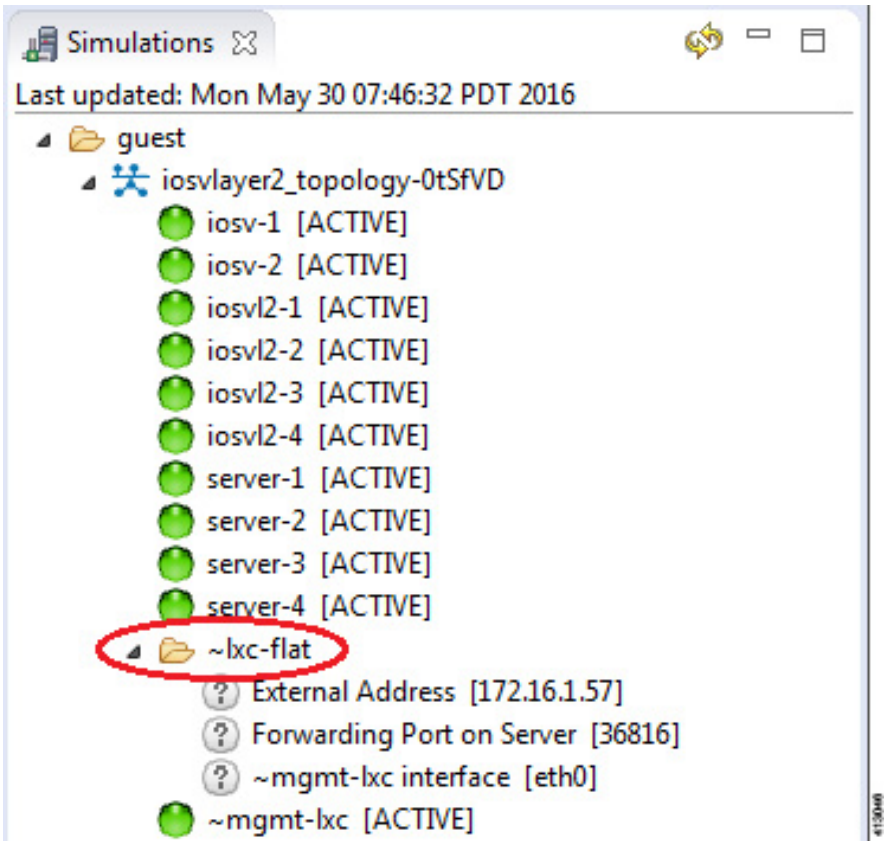

LXC facilitates SSH access to the VMs; it provides direct SSH access to each VM running inside the simulation. Telnet is not supported. As shown below, LXC is accessed by right-clicking the **~mgmt_lxc [ACTIVE]** node and selecting **SSH** from the list.

***Figure 12: Accessing LXC***

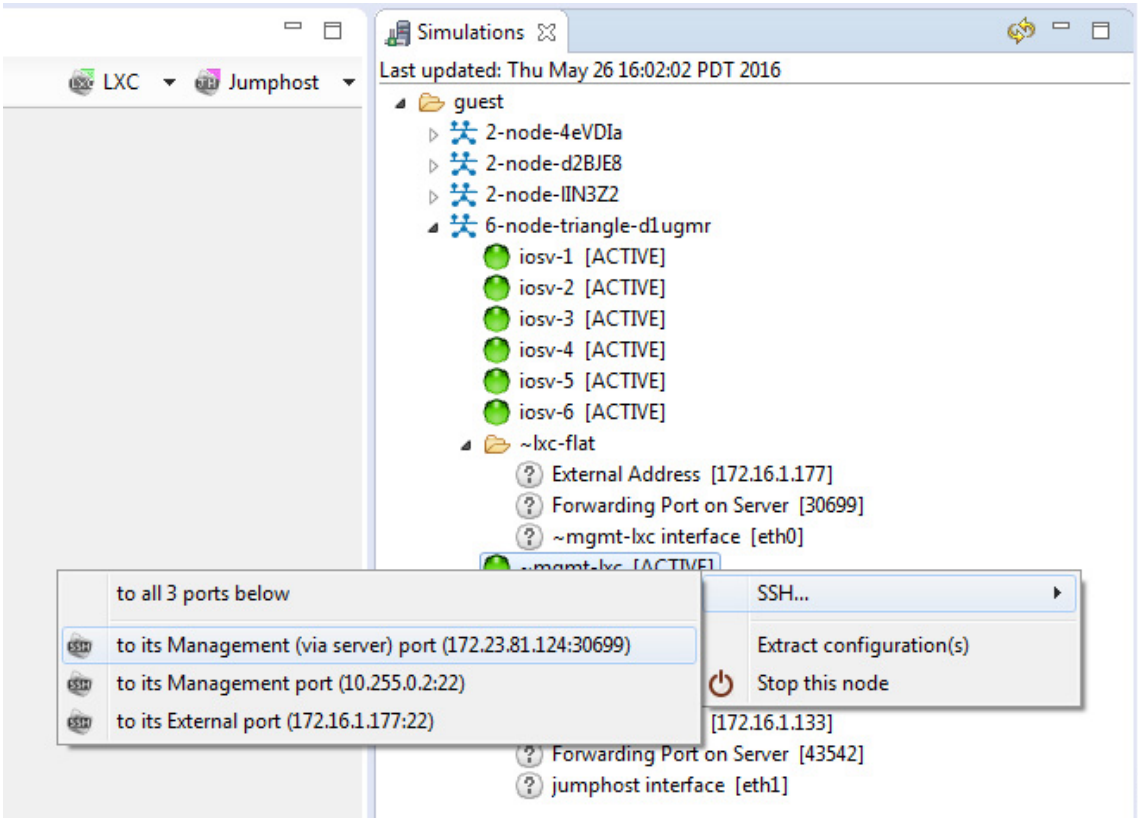


The LXC is automatically spun up and provides a jumphost point for access into the network. Connecting to the LXC means you can see the interfaces to the outside world and to the OOB network inside your simulation.

A connection is opened from the Cisco Modeling Labs client to a port on the Cisco Modeling Labs server and is forwarded to the LXC. The LXC, in turn, opens a connection to the Management Interface (Gi0/0) of the VM inside the simulation to the node instance.

## Static Port Assignment to the LXC

When the management network property **Private simulation network** is set, Cisco Modeling Labs assigns a random port for SSH port access to the LXC. However, you can statically define this by setting an extension on your topology.

To set an extension for your topology, complete the following steps.

**Step 1** Click on the canvas to open the **Topology** tab in the **Properties** view.

**Step 2** Click the **Extensions** tab.

**Step 3** Click the **Add new extension** icon.
The **Edit Extension** dialog box appears.

**Step 4** Enter the following values:

- **Key**: **lxc.host.tcp_port**
- **Value**: **30799**
- **Type**: **String** (from the drop-down list)

***Figure 13: Add a New Extension***

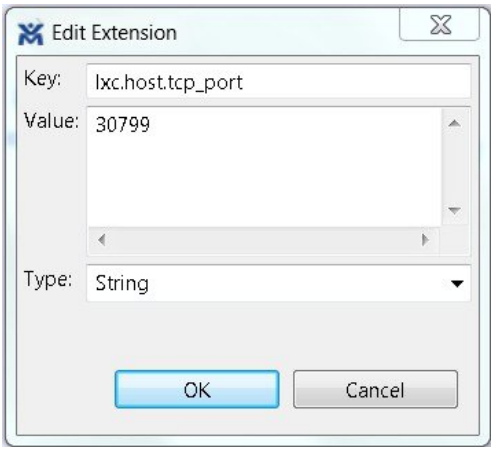


**Step 5** Click **OK** to add the new extension.
When the LXC starts, it will be bound to the TCP port specified in the new extension.

## LXC iPerf Container

The LXC iPerf container provides a stripped down lightweight Linux container which has been loaded with the iPerf application available from Downloads - iPerf.

iPerf is a tool for the active measurement of the maximum achievable bandwidth on IP networks. It supports tuning of various parameters related to timing, buffers and protocols (TCP, UDP, SCTP with IPv4 and IPv6). For each test it reports the bandwidth, loss, and other parameters.

***Figure 14: LXC iPerf Container***

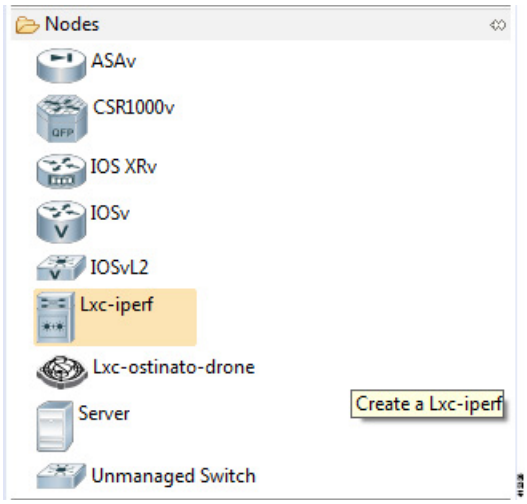


## LXC Ostinato Container

An LXC container is available that contains the Ostinato packet traffic generator application. This application provides data-plane traffic generation capabilities. The Ostinato **drone** (generator) is used in combination with the Ostinato GUI. The GUI can be obtained from Downloads – Ostinato.

When deployed, the LXC Ostinato container can be accessed using the SSH connection method.

**Note** Telnet does not work.

The Ostinato drone application executes automatically when the LXC becomes active.

**Note** The Ostinato drone application should not be installed on the host system as the version in the repositories cannot be executed in the LXC.

***Figure 15: LXC Ostinato Container***

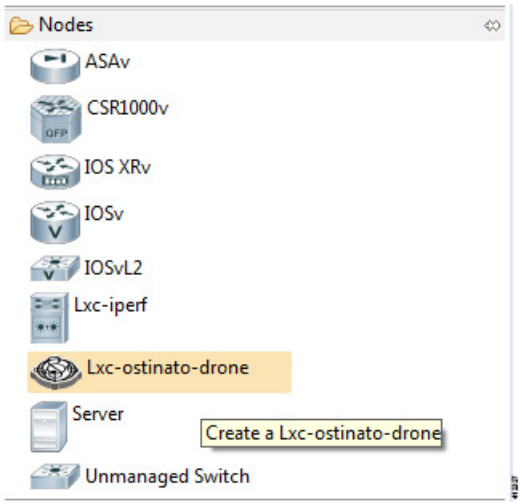


# Launch a Phased Simulation

On occasions, you may need to start your simulation in phases rather than having all nodes launched at the same time. This functionality is facilitated by the **Exclude node from simulation launch** check box, which allows you to pick and choose which nodes to start.

To launch a phased simulation, complete the following steps.

**Step 1** With the applicable topology open, click on the canvas and move the selection area over the nodes to be excluded from the running simulation. (Alternatively, you can double-click a specific node, hold down the **Shift** key, and select the remaining nodes.)

The **Properties** > **Node** view opens.

**Step 2** In the **Properties** > **Node** view, check the **Exclude node from simulation launch** check box.

***Figure 16: Nodes Excluded from Simulation Start***

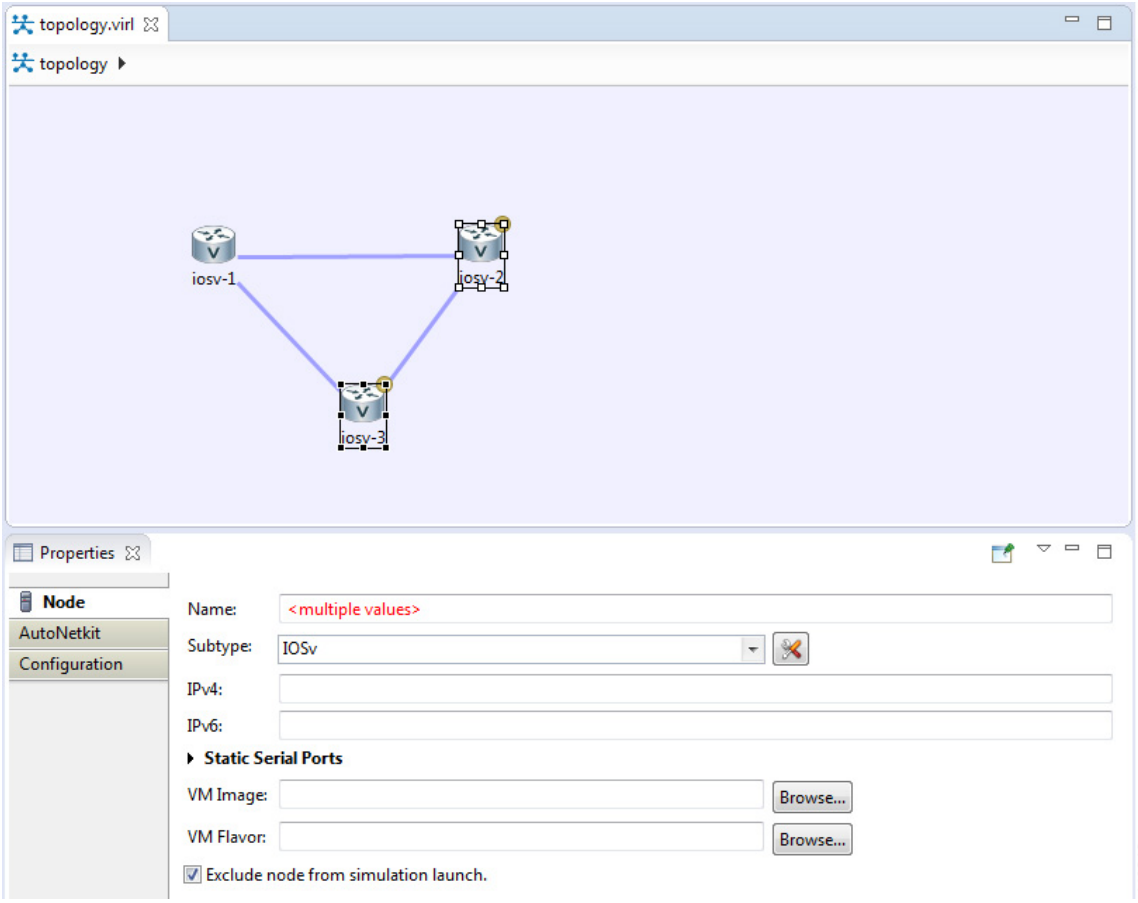


**Step 3** Save your topology using **Ctrl-S**.

**Note** You can set this property for each individual node, if you prefer. Select the node on the canvas, and select the **Exclude node from simulation launch** check box in the **Properties** > **Node** view.

**Step 4** From the toolbar, click the **Build Initial Configurations** button to build the node configurations.

**Step 5** From the toolbar, click the **Launch Simulation** button.
The simulation launches.
When the nodes are running on the Cisco Modeling Labs server, they are displayed in the **Simulations** view with the status as **[ACTIVE]**.

**Note** In the **Console** view, you can see the message **Node '<node name>' is excluded from automatic start** for the excluded nodes, and in the **Simulations** view, you can see that the excluded nodes have the state **[ABSENT]**.

***Figure 17: Phased Simulation Launched***

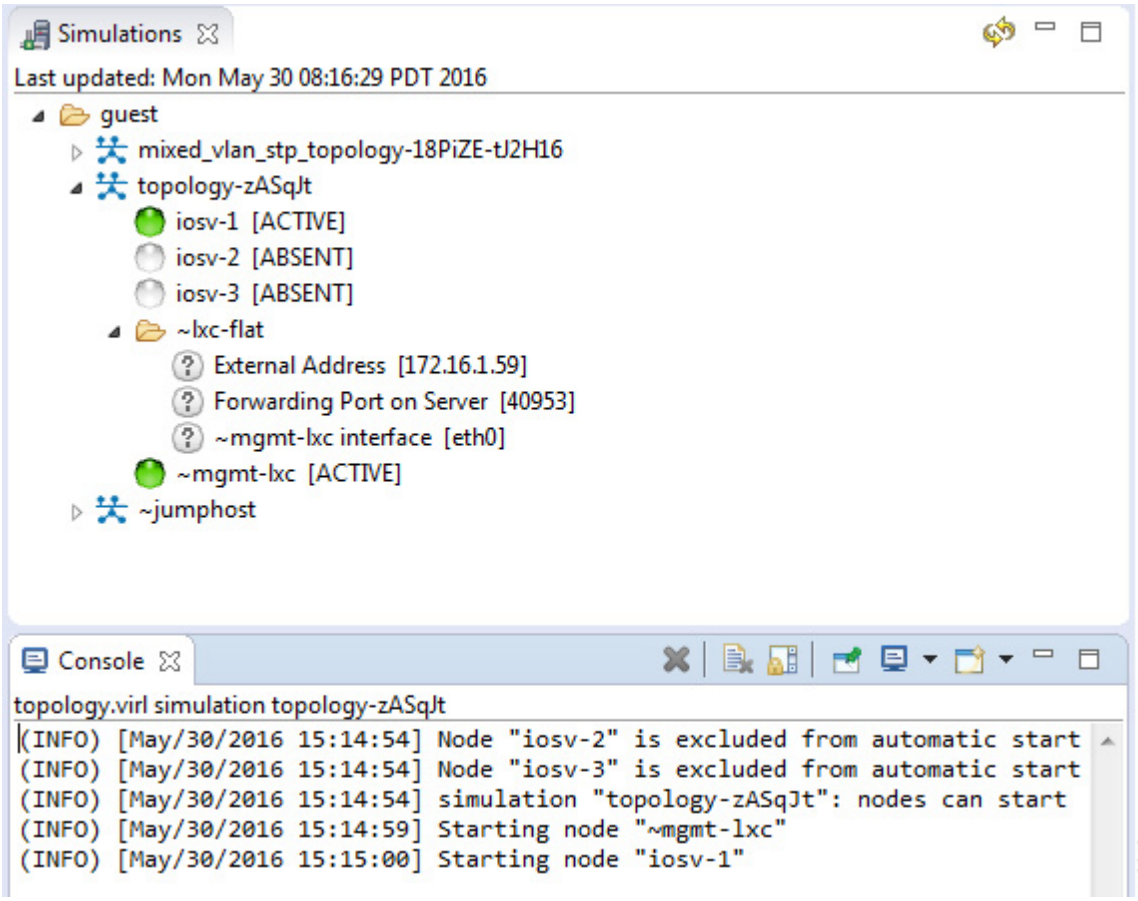

# Launch Simulation Options

In circumstances where you need to run a simulation for a specified time frame or you want to specify your own name for a simulation, complete the following steps:

**Step 1** From the menu bar, choose the **Simulation** button.

**Step 2** From the list, click **Launch Simulation with Options**.

***Figure 18: Launch Simulation with Options***

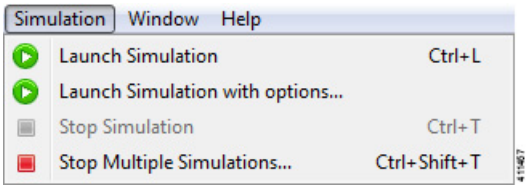


The **Launch Simulation** dialog box is displayed.

***Figure 19: Launch Simulation***

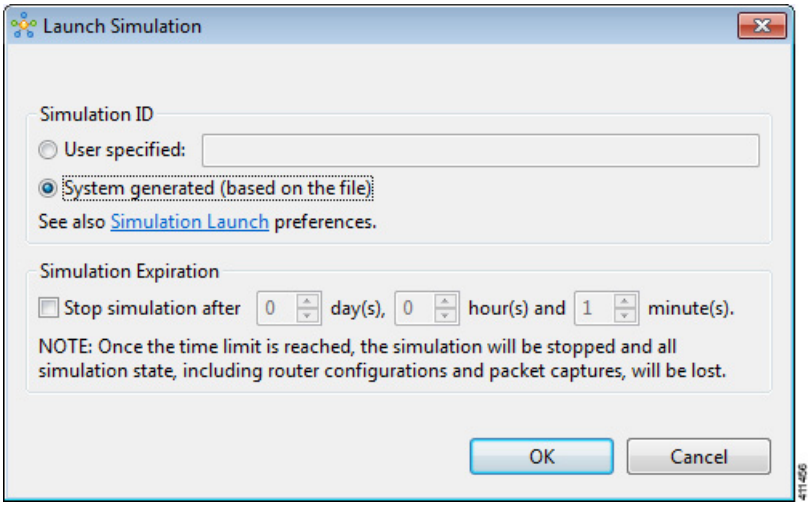


**Step 3** If you want to specify your own name or label for the simulation, check the **User Specified** radio button. Otherwise, leave the default **System generated ( based on a file)** radio button checked.

**Note** Alternatively, you can specify a simulation name via the **User Workspace Management** interface; select **My Simulations** > **Launch New Simulation** and enter the name in the **Simulation Name** field.

**Step 4** Set a time duration for the simulation by entering details for **Days**, **Hours**, and **Minutes** by using either the up and down arrows or entering the values directly. Click **OK** to apply your time limit to the simulation.
The simulation launches.

**Note** An expiration warning is displayed ten minutes (or less) before your simulation is due to expire.

***Figure 20: Simulation Expiration Warning***

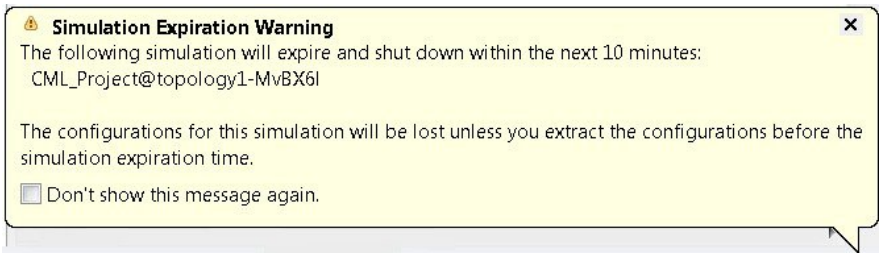


We recommend that you extract configurations for your simulation before it expires.

**Note** You can reset the time limit for a running simulation in the **User Workspace Management** interface. See Reset the Time Limit on a Running Simulation, on page 18 for more information.

## Reset the Time Limit on a Running Simulation

You can extend or reduce the time limit set for a running simulation in the **User Workspace Management** interface. To do this, complete the following steps.

**Step 1** Log in to the **User Workspace Management** interface.
On the **Overview** page, information on running simulations is displayed. Move to the applicable simulation under the **Session** heading. If your simulation is due to expire in ten minutes or less, the simulation name is displayed in red.

**Step 2** Under the **Options** column, click the down arrow and click **Set expiration**.

***Figure 21: Soon to Expire Simulation***

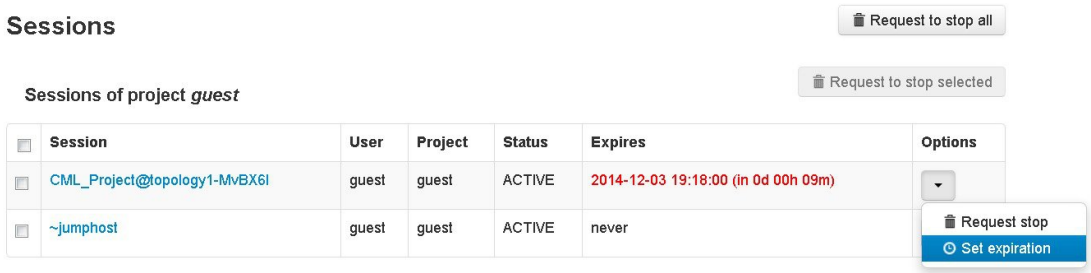

**Step 3** The **Set expiration for session** page is displayed. In the date and time field, enter either a new expiration date and time, a date only, or a time only for the simulation.

***Figure 22: Set Expiration for Session Page***

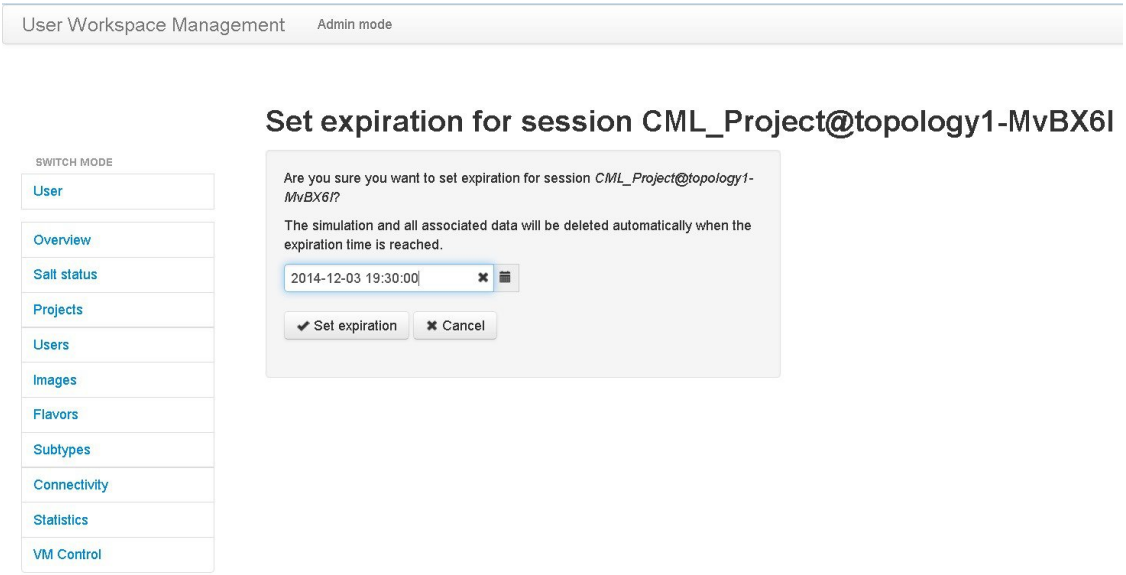


**Step 4** Click **Set expiration** to save the changes.
The time limit for the simulation is updated.

# Control Interface States

In a running simulation, you are able to change the state of the network interface by bringing it up or down.

**Note** This changes the state of the underlying communication infrastructure, not the interface state of the virtual machine.

To control the state of an interface, complete the following steps.

**Step 1** Log in to the **User Workspace Management** interface.
**Note** You must log in as a user other than the uwmadmin user, for example, guest.

**Step 2** On the **Overview** page, under **Sessions**, choose the applicable running session.

A list of active virtual machines and interfaces is displayed.

**Step 3** Scroll down to the **Interfaces** section and choose the applicable virtual machine.

**Step 4** From the applicable **Options** drop-down list, click **Update admin state**.

***Figure 23: Interface State Control Option***

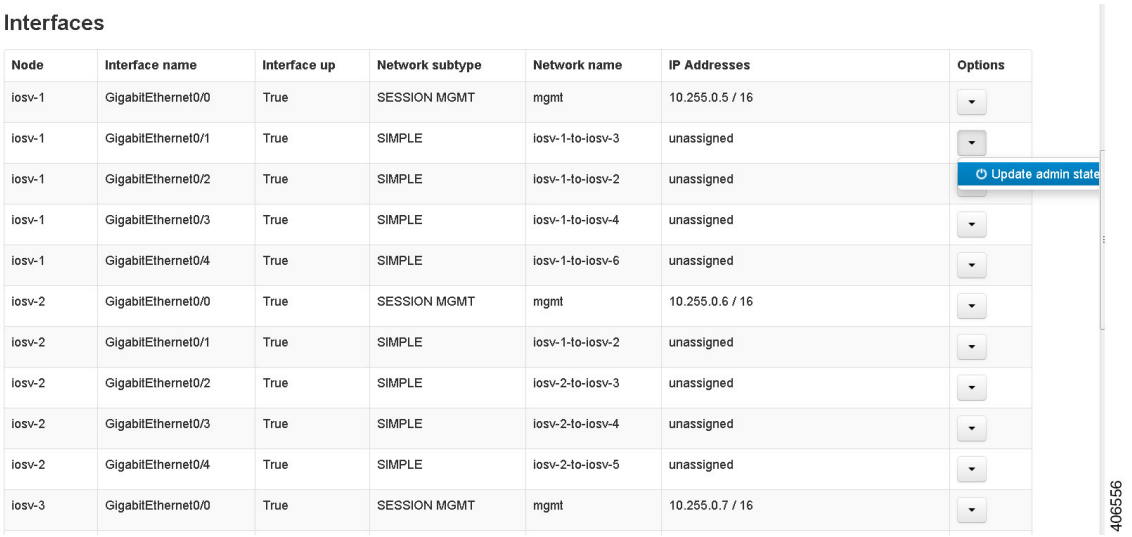


Depending on the current state of the interface, either a **Bring Down** or **Bring Up** page is displayed. In this case, the **Bring Down** page is displayed.

***Figure 24: Bring Down the Applicable Interface***

**Step 5** Click **Bring down** to bring down the network interface.

A message is displayed indicating that the interface has been brought down.

***Figure 25: Interface Successfully Brought Down***

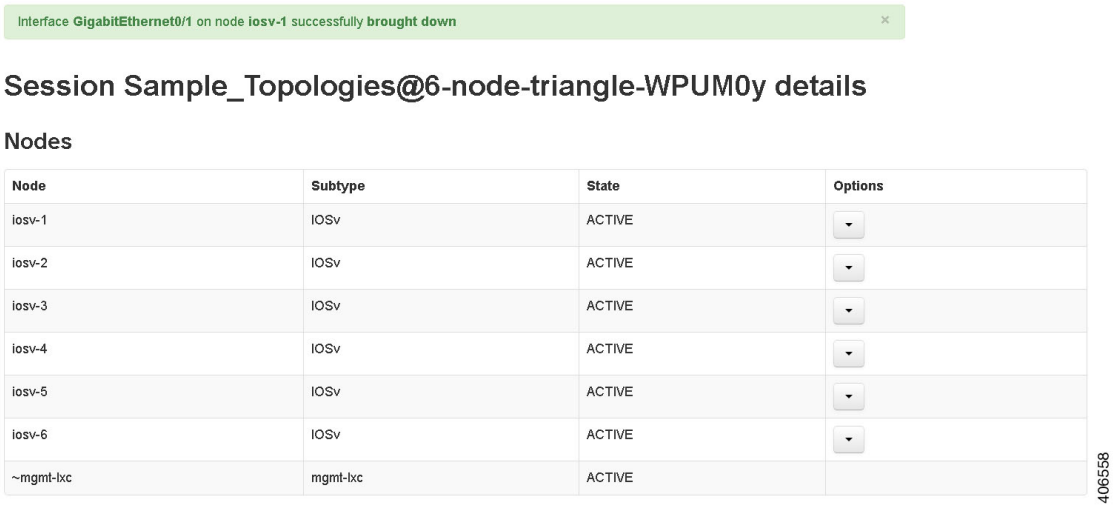


**Step 6** To bring up the interface at a later stage, click **Update admin state** again.
The **Bring up** page displayed.

***Figure 26: Bring Up the Applicable Interface***

**Step 7** Click **Bring up** to bring up the network interface.

A message is displayed indicating that the interface has been brought up.

***Figure 27: Interface Successfully Brought Up***

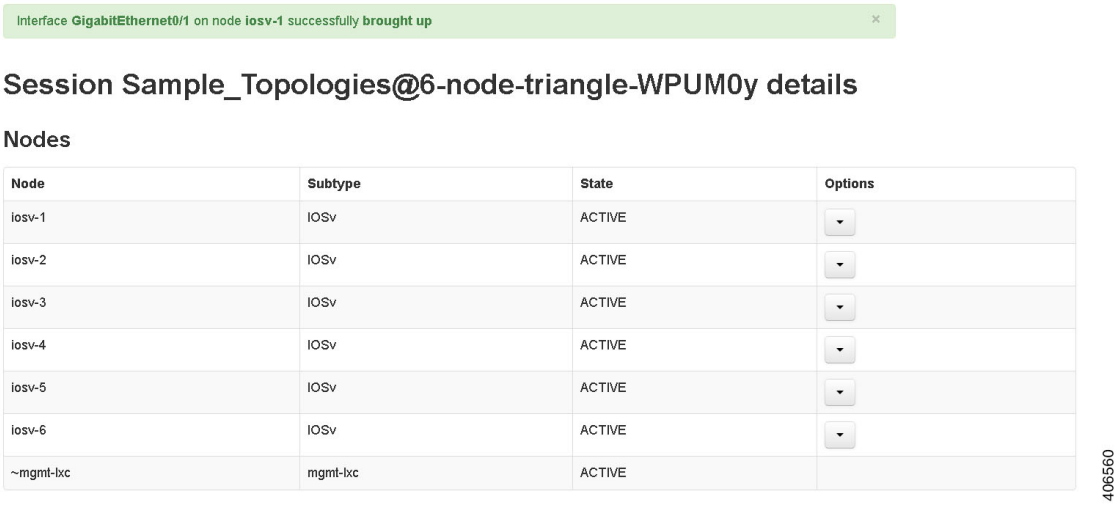


# Connect to a Simulation Node Console

Cisco Modeling Labs provides the capability for you to connect to your nodes via SSH and Telnet. You can start an SSH session, which connects into the node via the LXC, as described in Linux Container (LXC), on page 9.

This access method makes use of SSH to the LXC and then Telnet from the LXC to the nodes running inside the simulation. This does not use the console port of the nodes and is more reliable and faster to use.

# Connect to a Simulation Node Console via SSH

To connect to a simulation node console, complete the following steps.

### Before You Begin

- Launch a simulation.
- Ensure that the **Simulation** perspective is active.
- Ensure that the **Simulations** view and **Console** view are displayed.

**Step 1** To connect to a console for a specific node, right-click the node in the **Simulations** view and choose **SSH** > **to its Management (via LXC) port**.

*Figure 28: Connecting to a Node Console via SSH*

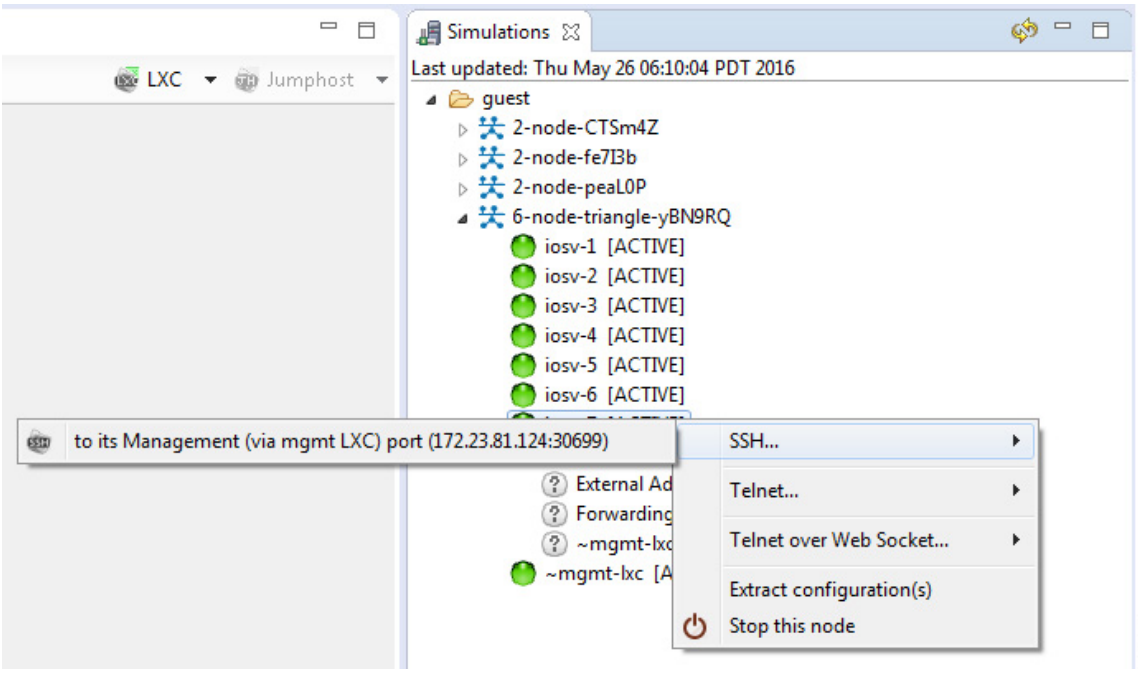

A new **Terminal** view opens.

***Figure 29: Terminal View***

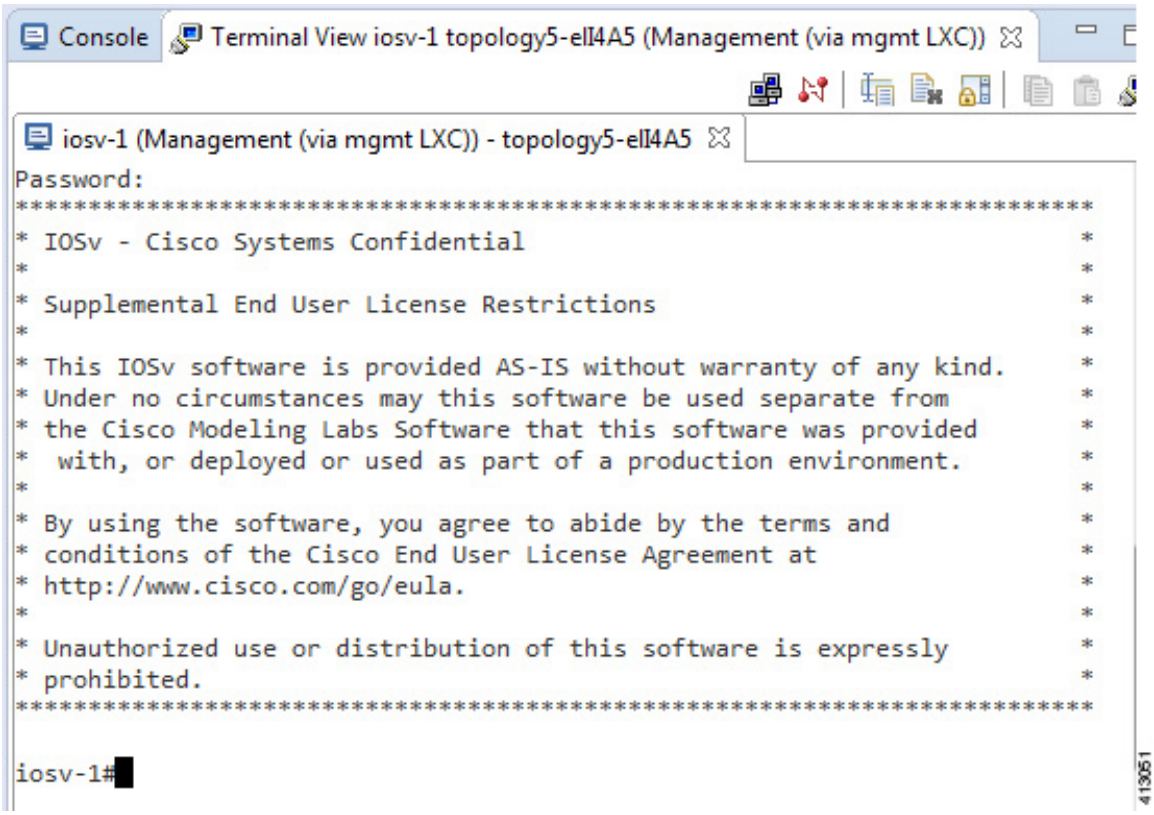


**Step 2** To disconnect a terminal from the simulation, click **Disconnect** in the **Terminal** view toolbar or click the **Close** icon in the **Terminal** view.

**Note** When you disconnect or close a **Terminal** view, all text in the view is discarded.

**Tip** When you disconnect a terminal but do not close the **Terminal** window, you can press **Enter** to reconnect the terminal.

## Connect to Multiple Simulation Node Consoles

To connect to all consoles for all nodes in a running simulation, complete the following steps.

**Step 1** Right-click the simulation in the **Simulations** view and choose **Telnet** > **to all <number> available Console ports**.

*Figure 30: Connect to Multiple Node Consoles*

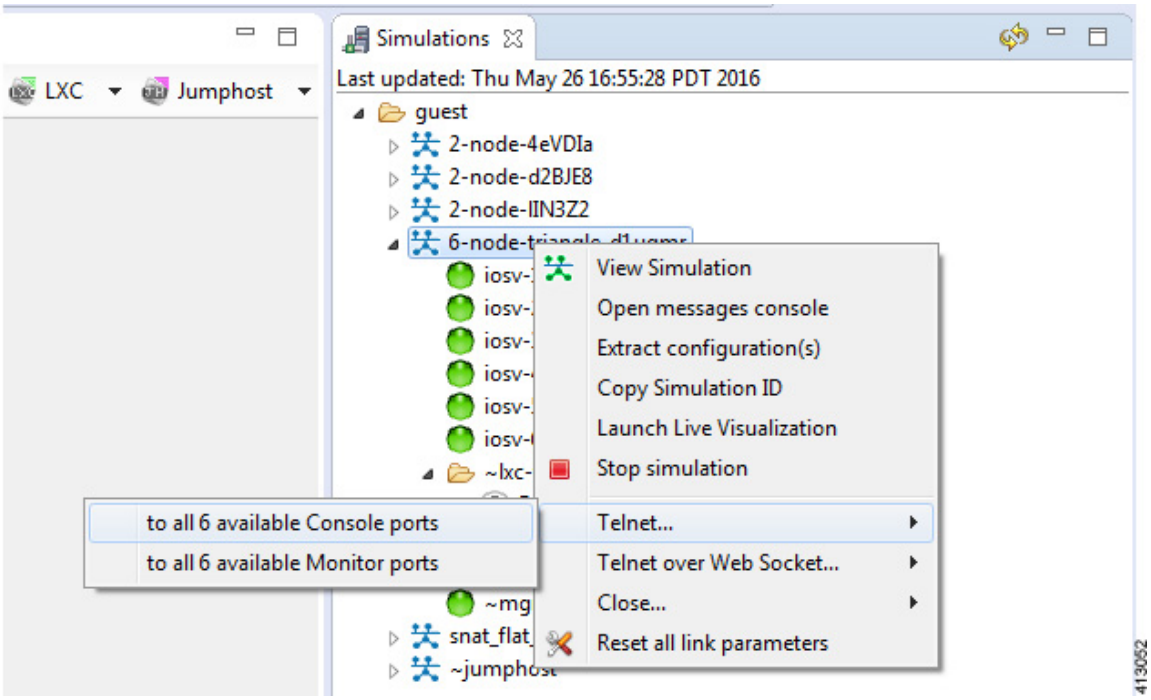

A new **Terminal** view opens for all console ports.

***Figure 31: Terminal View***

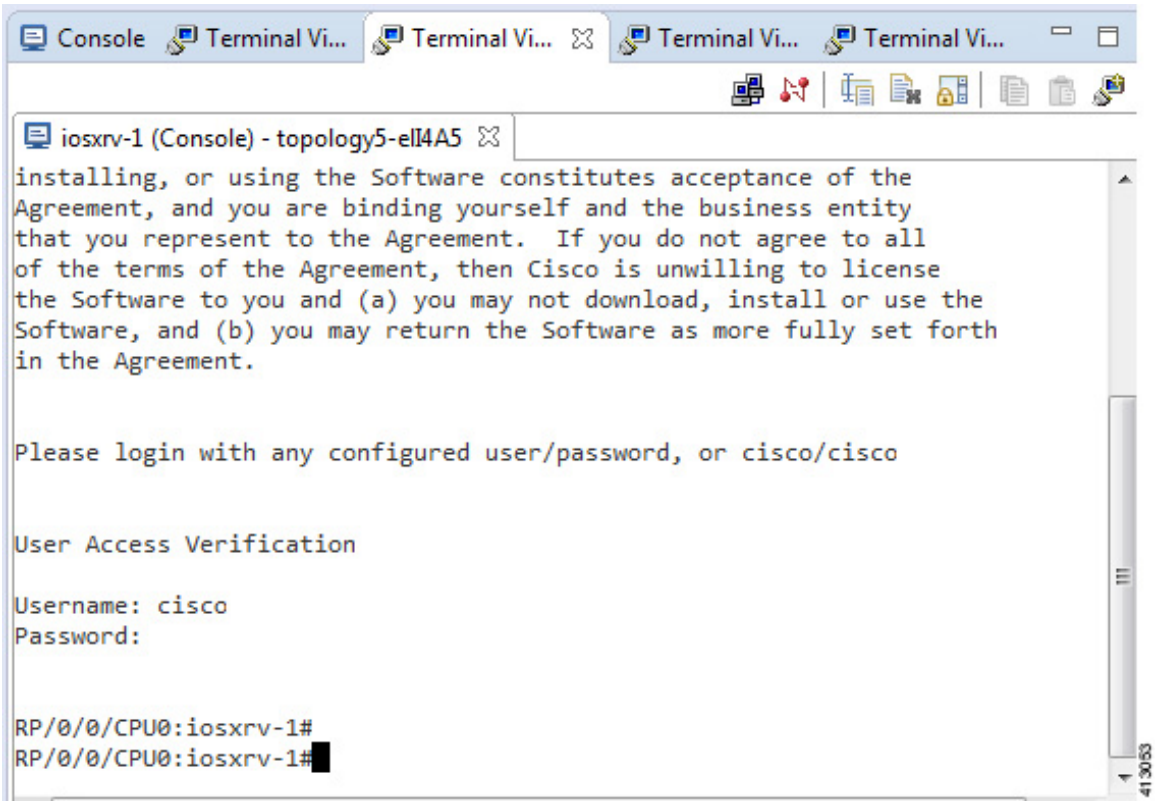


**Step 2** To disconnect a terminal from the simulation, click **Disconnect** in the **Terminal** view toolbar or click the **Close** icon in the **Terminal** view.

**Note** When you disconnect or close a **Terminal** view, all text in the view is discarded.

**Tip** When you disconnect a terminal but do not close the **Terminal** window, you can press **Enter** to reconnect the terminal.

## Terminal Multiplexer Functionality

A terminal multiplexer is available for use with the Cisco Modeling Labs client. It permits a number of terminals to be accessed and controlled from a single terminal. Terminals can be detached to run in the background and then reattached later.

The terminal multiplexer is available from **Window** > **Show View** > **Other** > **General** > **Terminal Multiplexer**.

When the terminal multiplexer starts up, a status line at the bottom of the terminal displays information on the current session and is used to enter interactive commands. It lists all of the active console sessions. You

can select console sessions individually or all together from the **View Menu** drop-down list. Keyboard commands entered in the command-line text box are broadcast to the selected sessions.

***Figure 32: Terminal Multiplexer***

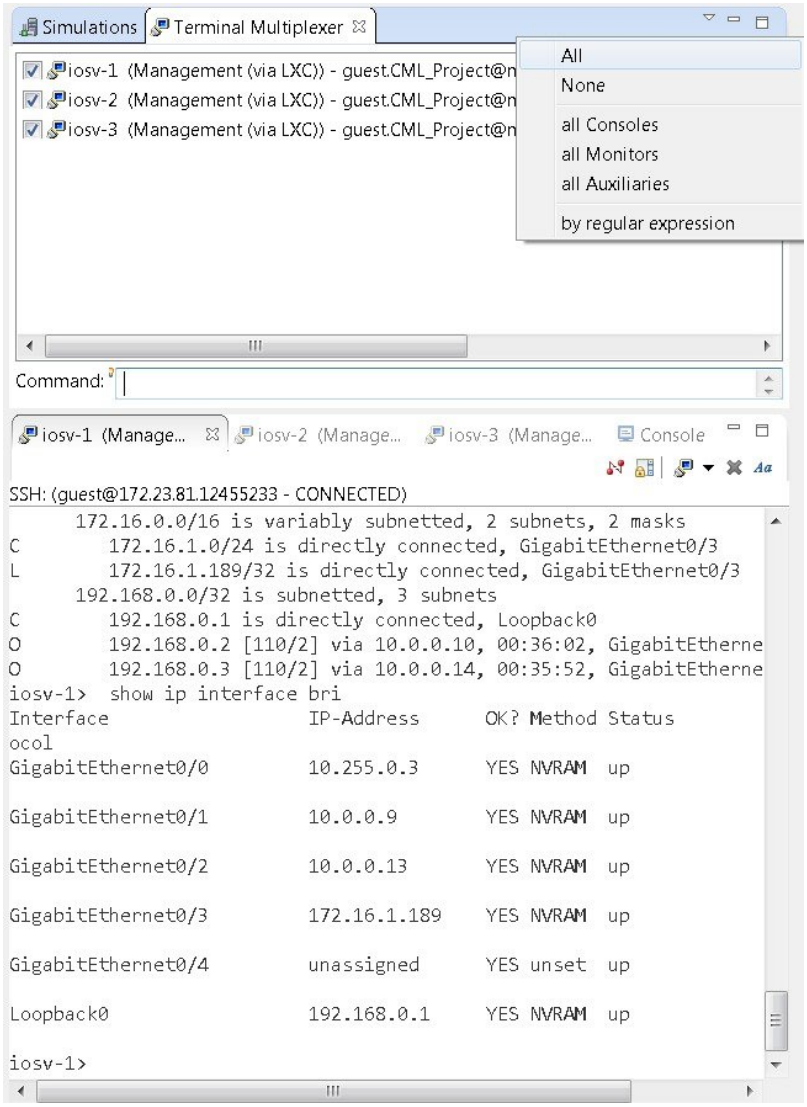

The terminal multiplexer also provides a command-line history, which you can access using **Ctrl-Space**.

***Figure 33: Accessing the Command-Line History***

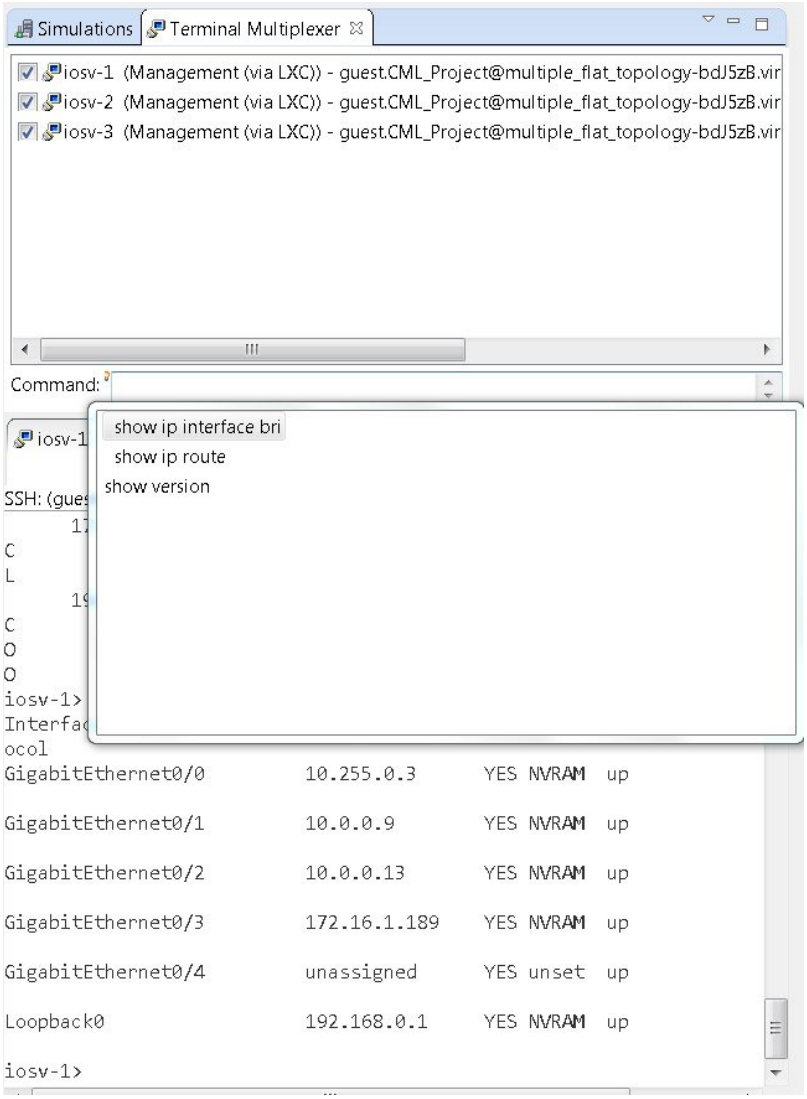


# Start a Single Node

To start a single node, complete the following steps.

**Step 1** Right-click a stopped node.

When a node is stopped, its status is shown as **[ABSENT]**.

**Step 2** Click **Start this node**.

***Figure 34: Start a Single Node***

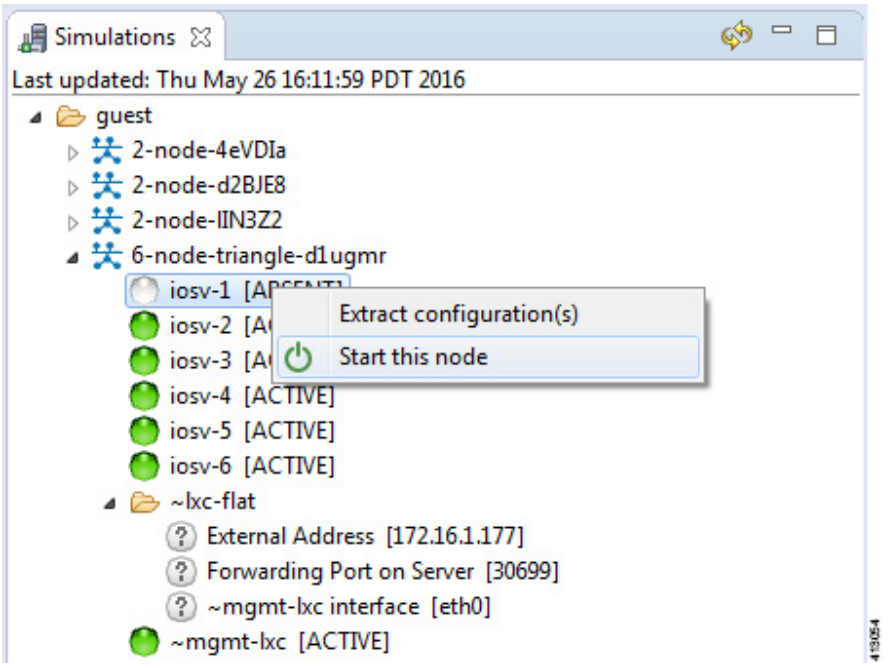


The **Requesting start** dialog box appears.

***Figure 35: Requesting Start Dialog Box***

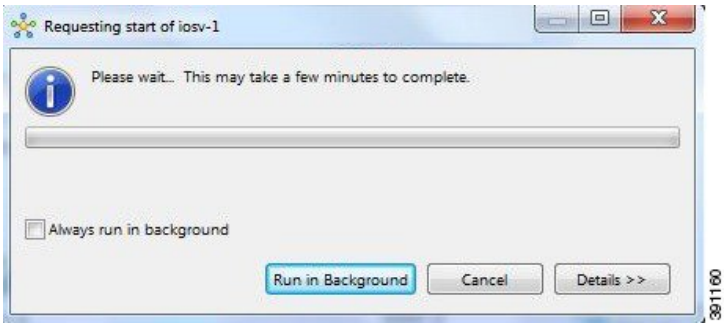


**Step 3** Choose one or more of the following actions:

a) Check the **Always run in the background** check box. All future node start requests, stop requests, and simulation launch requests run in the background and do not display dialog boxes.

**Note** To control the background setting, choose **File** > **Preferences** > **General**.

b) Click **Run in Background**.
The dialog box closes while the node simulation stops.

c) Click **Cancel** to return to the **Simulations** view.
d) Take no action and the node simulation restarts momentarily.

**Tip** When you click **Run in Background**, the status bar displays a progress icon. Click the icon to display a compact view of the progress. If an error is encountered, the icon displays a red **X**. Click the error icon to display the error dialog box.

# Start a Node in a Running Simulation

In cases where a phased simulation is running, you can later start those nodes not started with the initial simulation. To start a node in a running simulation, complete the following steps.

**Step 1** In the **Simulations** view, right-click the node.
**Note** A node not yet started has the status **[ABSENT]**.
**Step 2** Click **Start this node**.

***Figure 36: Start This Node Option***

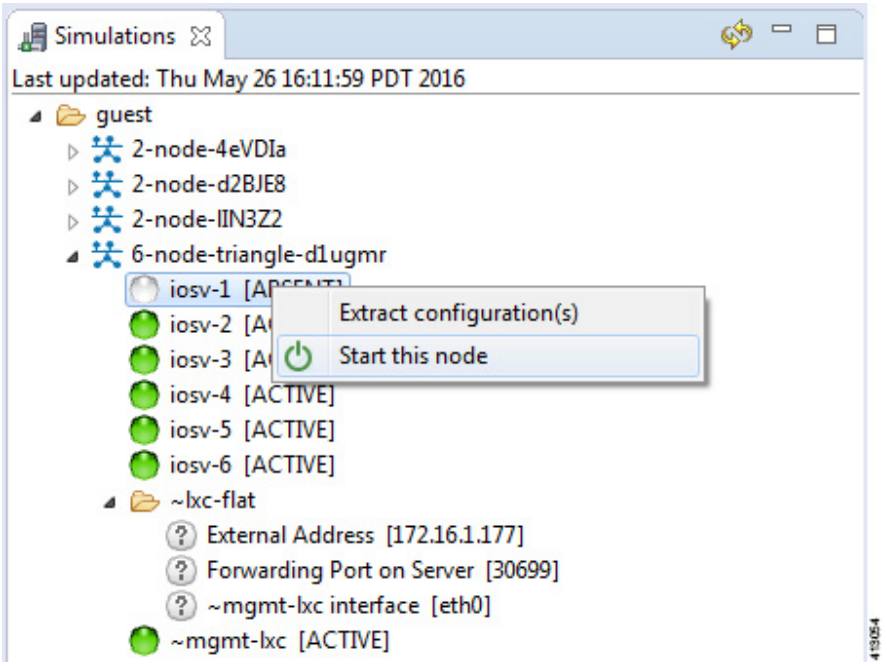


The **Requesting start** dialog box appears.

*Figure 37: Requesting Start Dialog Box*

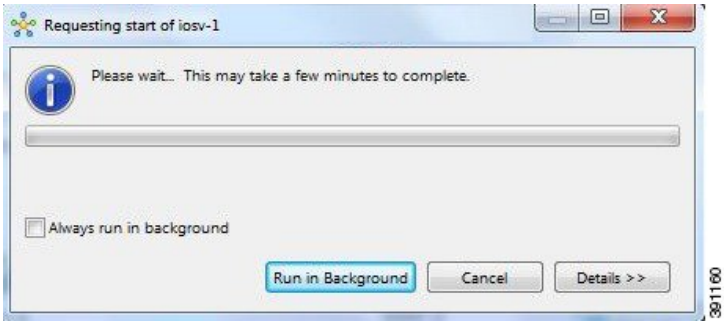


**Note** The status of the selected node is changed from **[ABSENT]** to **[ACTIVE]**, indicating that the node is up and running.

## Start Multiple Nodes in a Running Simulation

In cases where a phased simulation is running, you can later start those nodes not started with the initial simulation. To start multiple nodes, complete the following steps.

**Step 1** In the **Simulations** view, click the first node in the list to be started.

**Step 2** Hold down the **Shift** key and select the remaining nodes.

**Step 3** Right-click the selected nodes.

**Step 4** Click **Start nodes**.

***Figure 38: Start Nodes Option***

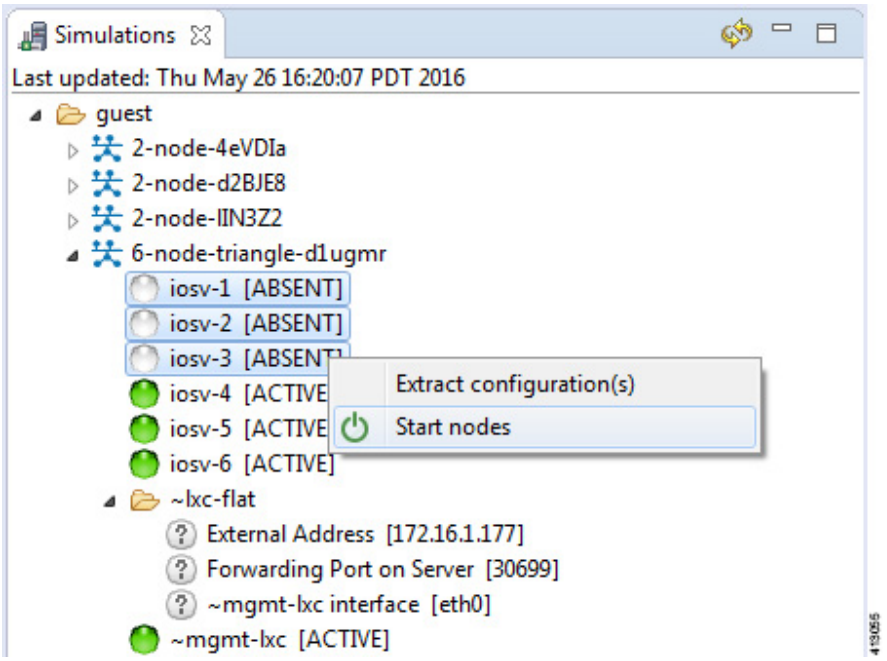


The **Requesting start** dialog box appears.

***Figure 39: Requesting Start Dialog Box***

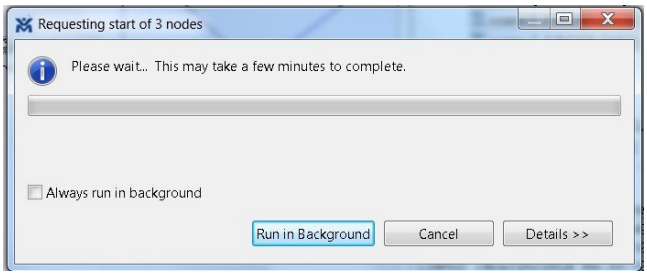

**Note** The status of the selected nodes changes from **[ABSENT]** to **[ACTIVE]**, indicating that the nodes are up and running.

***Figure 40: Multiple Nodes Started***

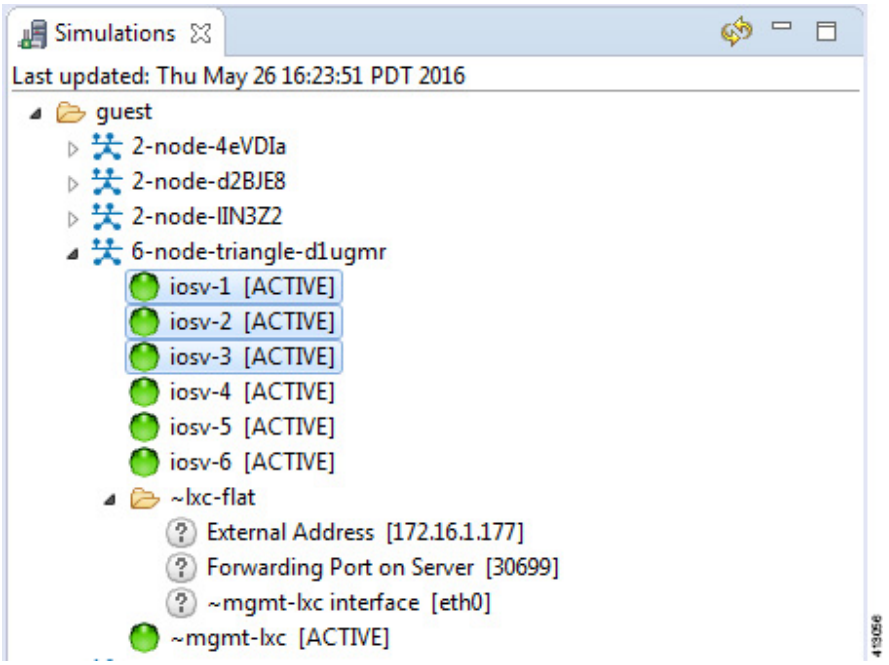


# Stop a Simulation

There are several ways to stop a simulation. In addition, you can stop multiple simulations at the same time. These are discussed in the following sections.

## Stop a Simulation from the Toolbar

To stop a simulation from the toolbar, complete the following steps.

**Step 1** In the toolbar, click the **Stop Simulations** button.

A **Stop Simulation(s)** dialog box appears.

***Figure 41: Stop Simulation(s) Dialog Box***

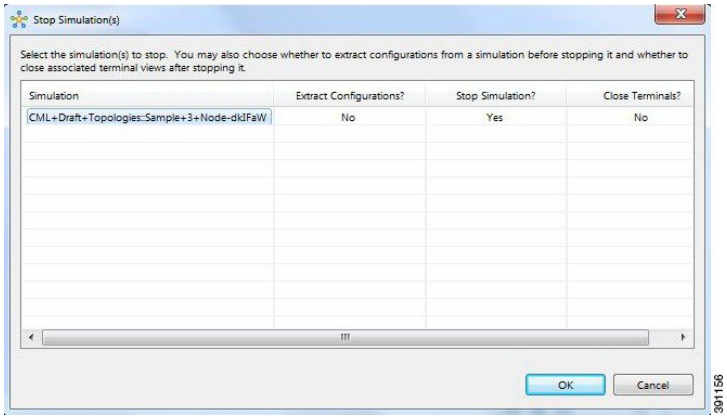


**Step 2** In the **Simulation** column, click once to highlight the simulation to stop.

**Step 3** (Optional) To save the configurations, click the adjacent setting in the **Extract Configurations?** column until the prompt changes to **Yes**.

**Note** Configurations for server nodes are not extracted.

**Step 4** (Optional) To close the internal terminals associated with the simulation, click the adjacent setting in the **Close Terminals?** column until the prompt changes to **Yes**.

**Note** External terminal connections are not stopped as part of this operation and must be closed manually.

**Step 5** To stop the simulation, click the adjacent setting in the **Stop Simulation?** column until the prompt changes to **Yes**.

**Step 6** Click **OK** to stop the simulation, or click **Cancel** to leave the simulation running.

On OS X, you update the values for **Extract Configurations?**, **Stop Simulation?**, and **Close Terminals?** in the columns directly. You do not need to select the name of the simulation.

# Stop a Simulation from the Simulations View

To stop a simulation, complete the following steps.

**Step 1** In the **Simulations** view, right-click the simulation name and select **Stop Simulation**.

***Figure 42: Stop a Simulation***

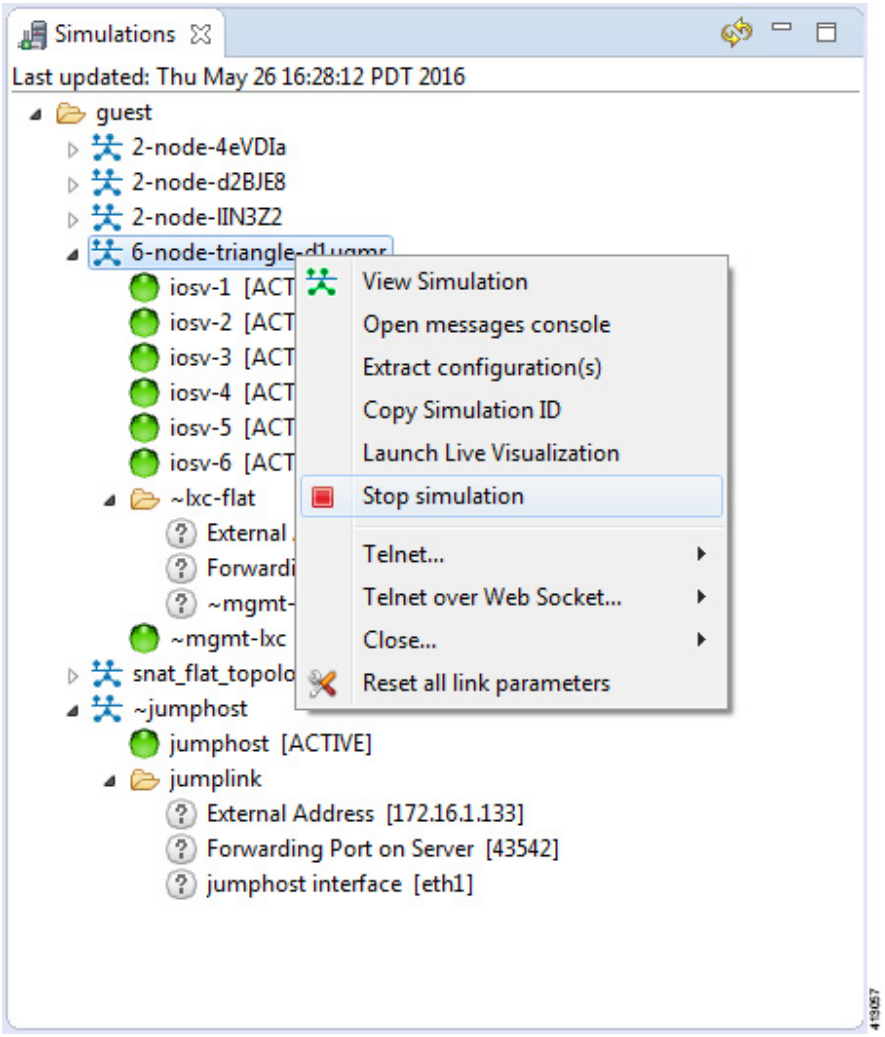

The **Stop Simulation Confirmation** dialog box appears.

***Figure 43: Stop Simulation Confirmation Dialog Box***

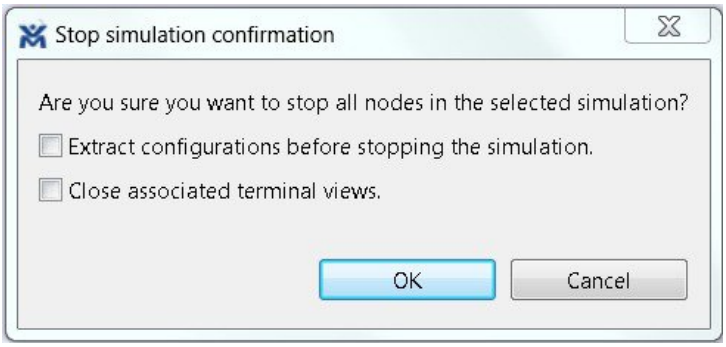


- (Optional) Check the **Extract Configurations before Stopping the Simulation** check box to save the current configurations.
- (Optional) Check the **Close Associated Terminal Views** check box to close all the open internal terminals associated with the specific simulation. External terminals are not closed.

**Step 2** Click **OK** to stop the simulation.

Once selected, all nodes in the simulation start shutting down. It may take a few minutes for the simulation to shut down completely and to disappear from the **Simulations** view.

**Note** For instances where a user account expires, all running simulations for that user continue to run. Since the user account can no longer log in to stop them, they will remain active until the next system reboot or until the system administrator explicitly stops them.

# Stop Multiple Simulations from the Simulations View

To stop multiple simulations, complete the following steps.

**Step 1** In the **Simulations** view, click the first simulation in the list to stop.

**Step 2** Hold down the **Shift** key and select the remaining simulations.

**Step 3** Right-click the selected simulations and select **Stop Simulations**.

***Figure 44: Stop Multiple Simulations***

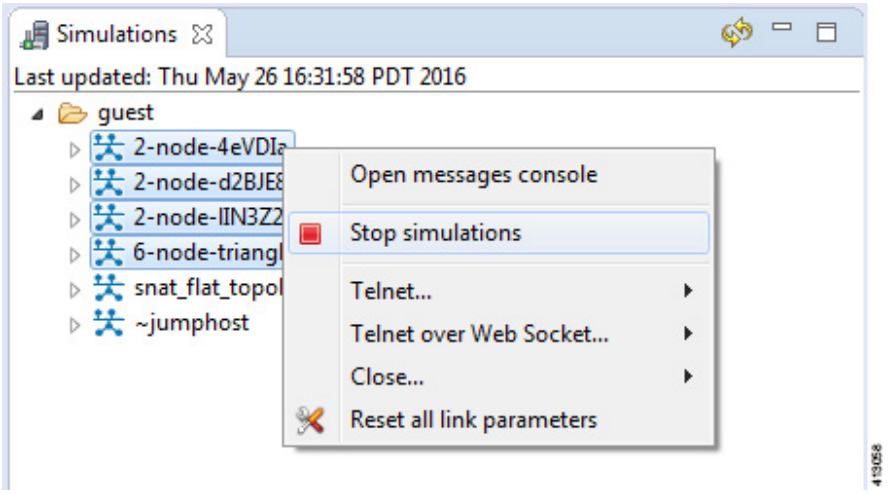


The **Stop Simulations Confirmation** dialog box appears.

***Figure 45: Stop Simulations Confirmation Dialog Box***

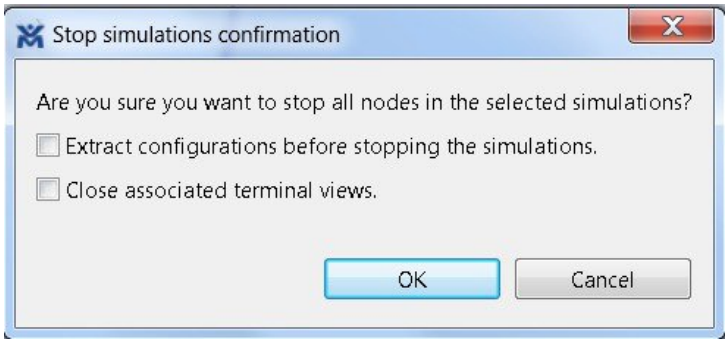


- (Optional) Check the **Extract Configurations before Stopping the Simulation** check box to save the current configurations.
- (Optional) Check the **Close Associated Terminal Views** check box to close all the open internal terminals associated with the specific simulation. External terminals are not closed.

**Step 4** Click **OK** to stop the simulations.
Once selected, all nodes in the simulations start shutting down. It may take a few minutes for the simulations to shut down completely and to disappear from the **Simulations** view.

**Note** For instances where a user account expires, all running simulations for that user continue to run. Since the user account can no longer log in to stop them, they will remain active until the next system reboot or until the system administrator explicitly stops them.

# Stop a Single Node

To stop a single node in a simulation, complete the following steps.

**Step 1** In the **Simulations** view, right-click the node to stop and select **Stop this Node**.

***Figure 46: Stop a Single Node***

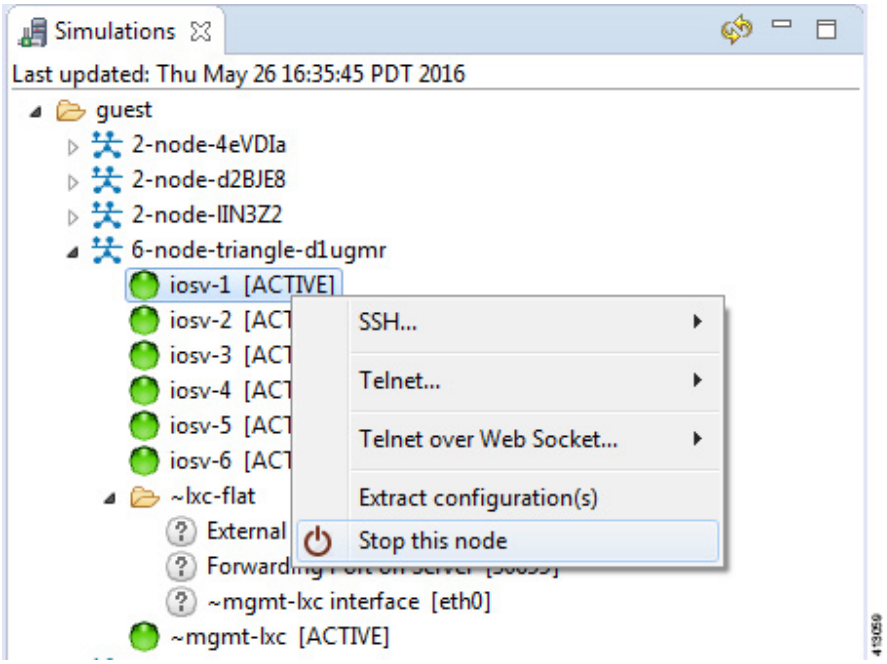


The **Are you sure?** dialog box appears.

**Step 2** Click **OK** to stop the node. Alternatively, click **Cancel** to abandon the operation and return to the simulation.

**Note** When you click **OK**, the node stops without saving any changes to the configuration.

When a node is stopped, its status changes to **[ABSENT]**.

# Stop Multiple Nodes

To stop multiple nodes in a running simulation, complete the following steps.

**Step 1** In the **Simulations** view, click the first node in the list to stop.

**Step 2** Hold down the **Shift** key and select the remaining nodes.

**Step 3** Right-click the selected nodes.

**Step 4** Click **Stop Nodes**.

***Figure 47: Stop Multiple Nodes***

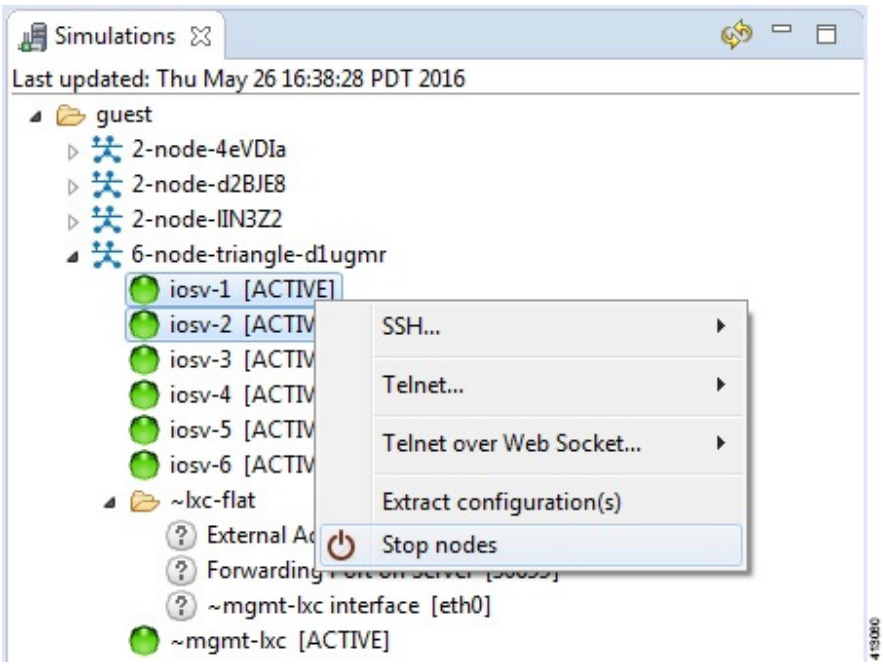


The **Are you sure?** dialog box appears.

**Step 5** Click **OK** to stop the nodes. Alternatively, click **Cancel** to abandon the operation and return to the simulation.

**Note** When you click **OK**, the nodes stop without saving any changes to the configuration.

When the nodes are stopped, their status changes to **[ABSENT]**.

***Figure 48: Multiple Nodes Stopped***

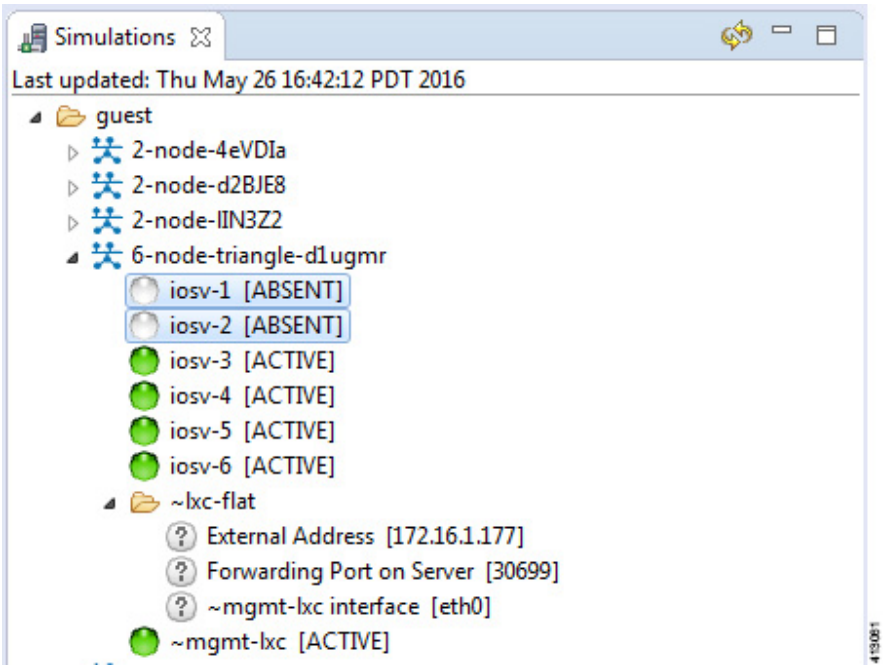


# Modify a Node Configuration in the Simulation

You can modify node configurations in a running simulation. To do this, the type of connection available, either SSH or Telnet depends on the option set for the topology property **Management Network**.

- When **Management Network** is set to **Private simulation network**, SSH and Telnet are available.
- When **Management Network** is set to **Shared flat network**, **Private project network**, or **not specified**, only Telnet is available.

## Modify a Node Configuration in the Simulation via SSH

To modify a node configuration in a running simulation via SSH, complete the following steps.

**Step 1** Right-click the node in the **Simulations** view and choose **SSH > to its Management (via LXC) port**.

*Figure 49: Connecting to a Node Console*

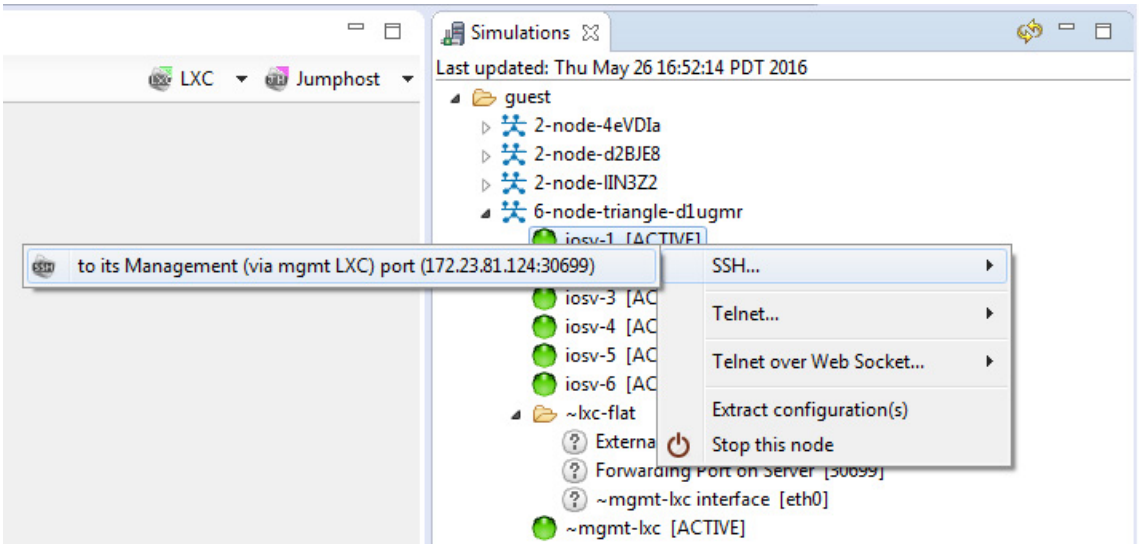

A new **Terminal** view opens.

***Figure 50: Terminal View***

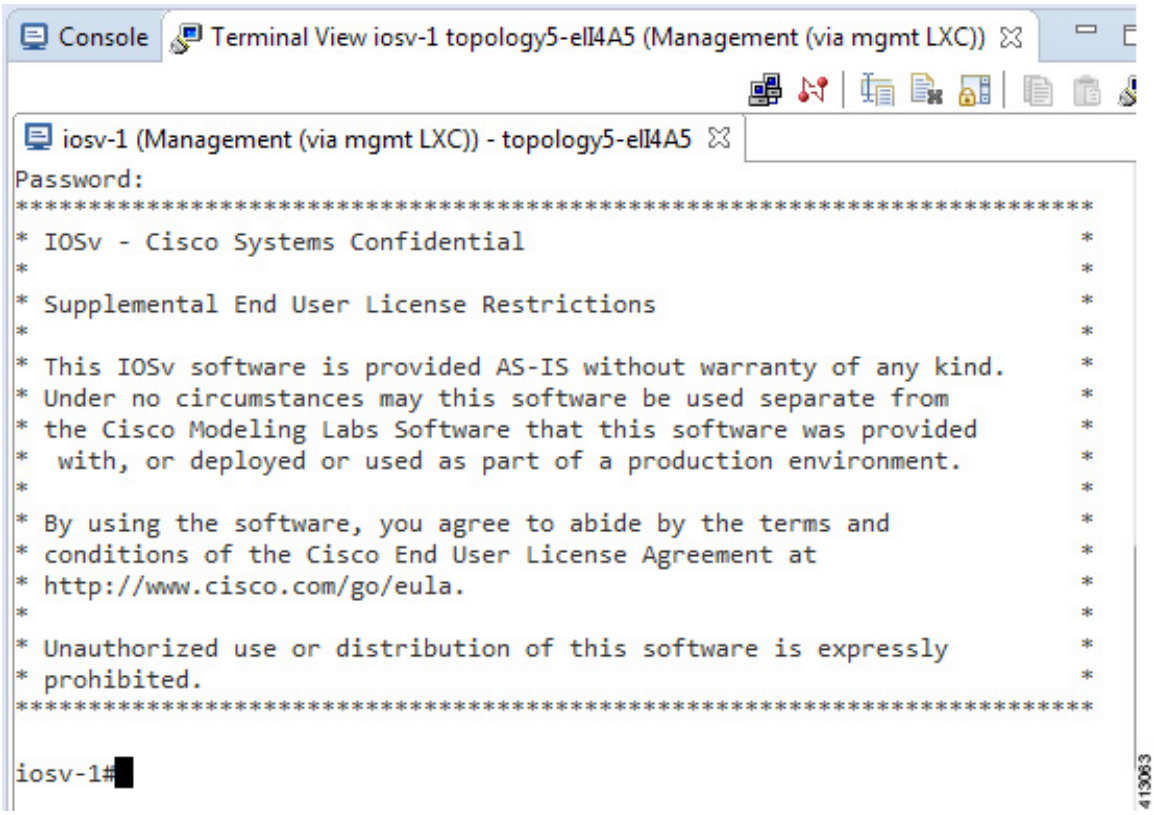


**Step 2** If no banner or router prompt is visible, press **Enter**.
You are now working with the operating system running on the node, for example, Cisco IOSv virtual software.

**Step 3** Use the operating system commands to view or modify the node configuration.

**Note** Changes you make to the configuration do not appear in the canvas of the **Topology Editor**.

# Modify a Node Configuration in the Simulation via Telnet

To modify a node configuration in a running simulation via Telnet, complete the following steps.

**Step 1** Right-click the node in the **Simulations** view and choose **Telnet > to its Console port**.

*Figure 51: Connecting to a Node Console*

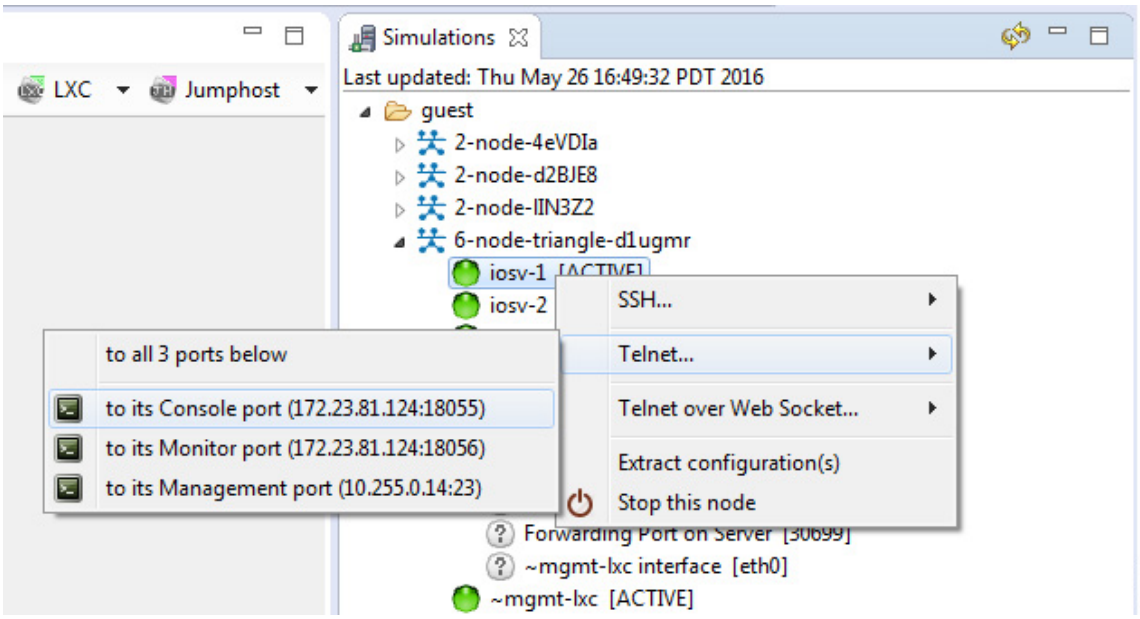

A new **Terminal** view opens.

***Figure 52: Terminal View***

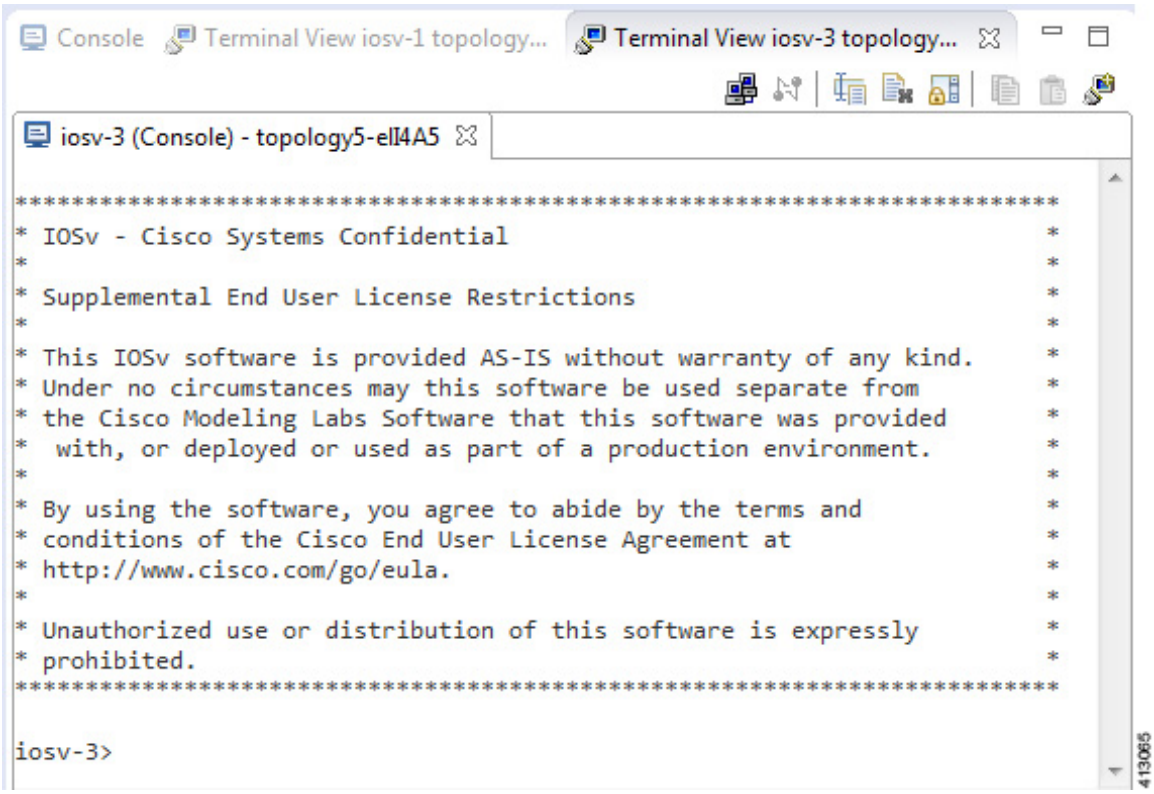


**Step 2** If no banner or router prompt is visible, press **Enter**.
You are now working with the operating system running on the node, for example, Cisco IOSv virtual software.

**Step 3** Use the operating system commands to view or modify the node configuration.

**Note** Changes you make to the configuration do not appear in the canvas of the **Topology Editor**.

# Modify Multiple Node Configurations in the Simulation

To modify multiple node configurations in a running simulation, complete the following steps.

**Step 1** Right-click the topology in the **Simulations** view and choose **Telnet > to all <number> available Console ports**.

*Figure 53: Connect to Multiple Node Consoles*

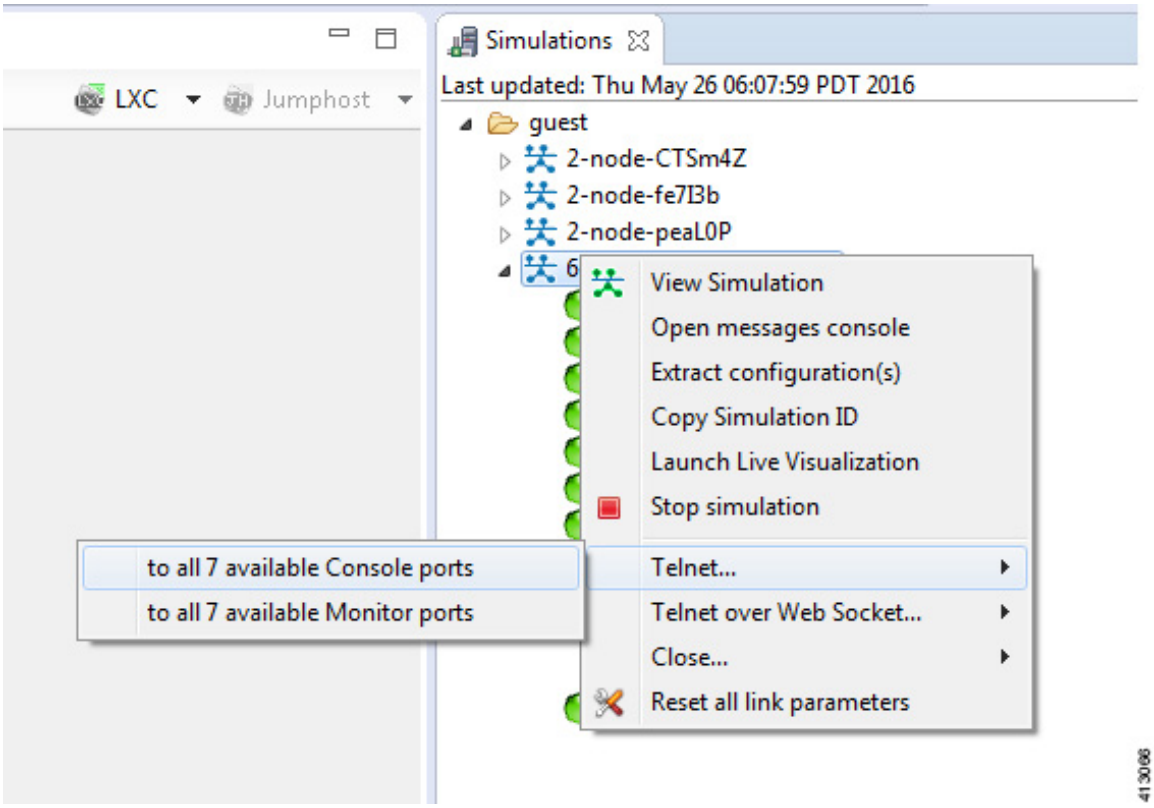

A new **Terminal** view opens for each of the consoles.

***Figure 54: Terminal Views***

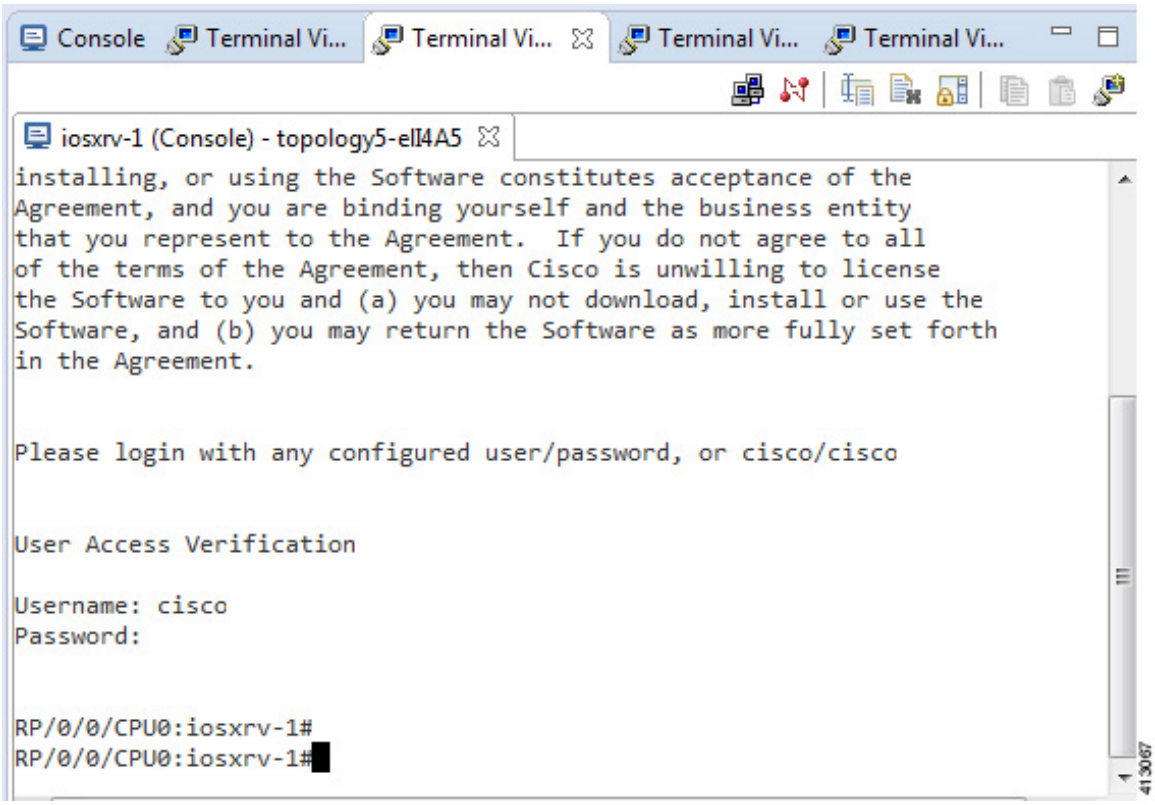


**Step 2** If no banner or router prompt is visible, press **Enter**.
You are now working with the operating system running on the node, for example, Cisco IOSv virtual software.

**Step 3** Use the operating system commands to view or modify the node configuration.
**Note** Changes you make to the configuration do not appear in the canvas of the **Topology Editor**.

# Extract and Save Modified Configurations

To extract and save modified configurations, complete the following steps.

### Before You Begin

- You have modified a configuration within one or more nodes running within the simulation and want to save the changes.

- Ensure that all routers in the simulation are operational before attempting to extract their configurations.

**Step 1** In the **Simulations** view, right-click the topology name, making sure not to click the node name, and select **Extract Configurations**.

***Figure 55: Simulations View Context Menu***

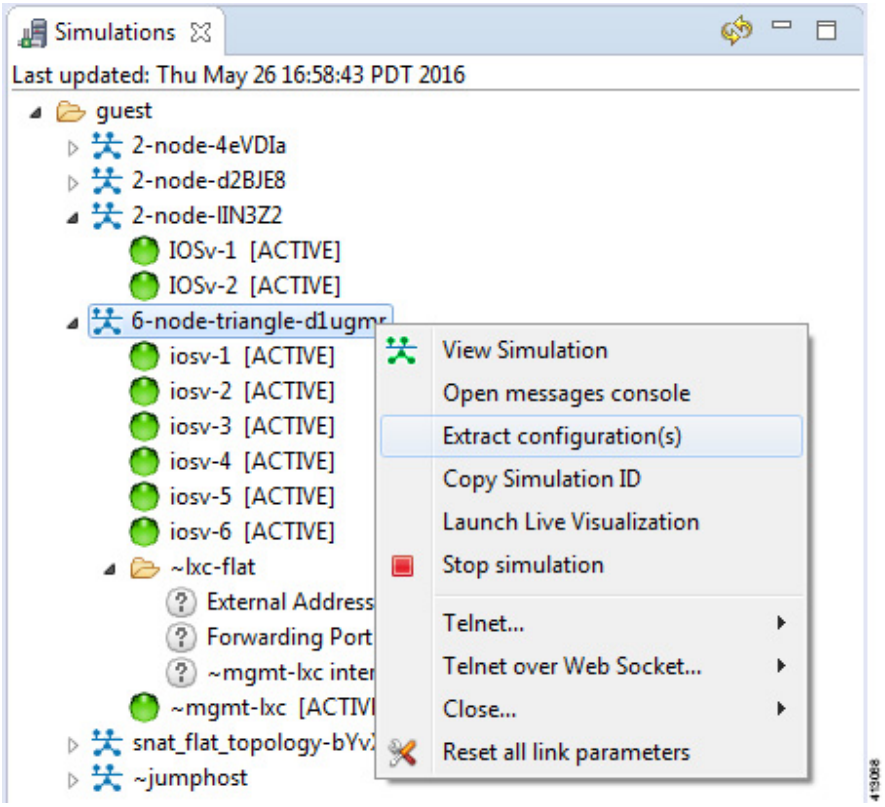


A confirmation dialog box appears.

**Step 2** Ensure that all external Telnet connections to the simulation are closed before proceeding.

**Note** You must close all external Telnet connections to the simulation before you can proceed.

***Figure 56: Extract Configurations Dialog Box***

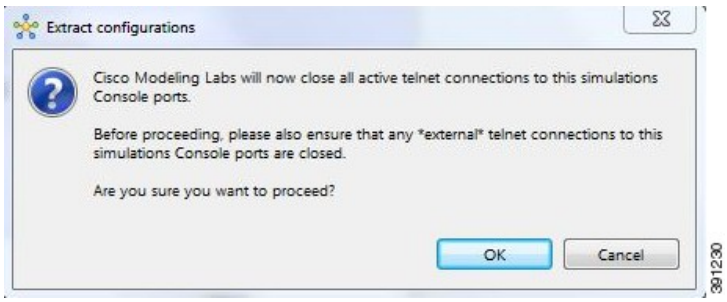


**Step 3** Click **OK**.
The Extracting Configurations dialog box appears indication that the extraction is in process.

When the extraction is complete, a message is displayed.

***Figure 57: Extraction Complete Message***

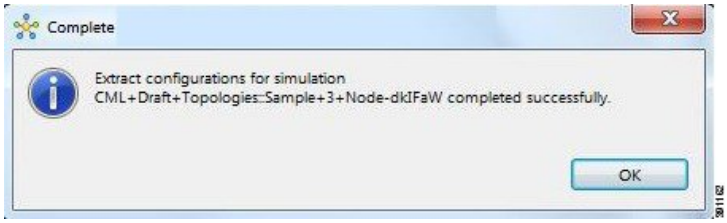


**Note** The configuration information is extracted and saved in the filename.virl file that corresponds to the topology. For example, the file New_Topology that is running as a New_Topology-dkIFaW simulation has its configuration extracted to the file **New_Topology.virl**.

**Step 4** Click **OK**.

# Partial Configuration Extraction

During a configuration extraction, if the process encounters issues or fails for a particular node, the problem node is identified and reported.

The extraction process then continues for all other nodes in the simulation and returns collected configurations to you.

# Linux Server Snapshot Support

When a Linux server is present in a running simulation, you can use the **User Workspace Management** interface to take a snapshot of the disk content of the server. This newly created user-specific disk image can be used in other simulated sessions.

To take a snapshot of the server's disk contents, complete the following steps.

**Step 1** Log in to the **User Workspace Management** interface.

**Note** You must log in as a user other than the uwmadmin user, for example, guest.

**Step 2** On the **Overview** page, under **Sessions**, select the applicable running simulation.

***Figure 58: Running Simulation Listed***

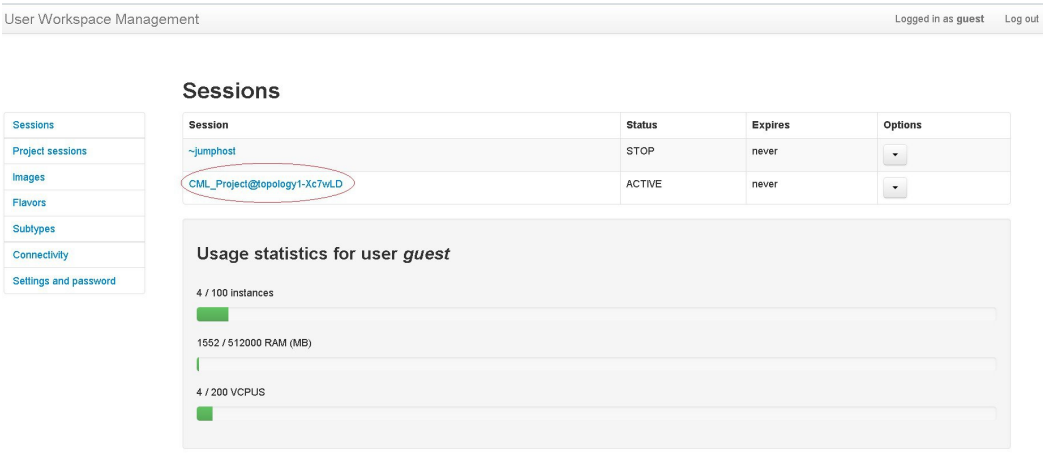


A list of active VMs is displayed.

***Figure 59: Active VMs***

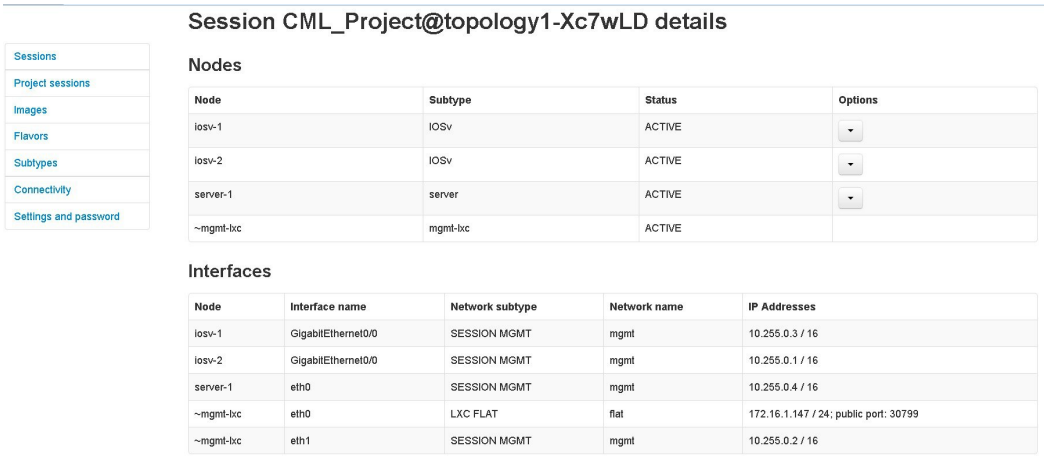

**Step 3** Select the applicable Linux server, and from the **Options** drop-down menu, click **Create snapshot**.

***Figure 60: Create Snapshot Option***

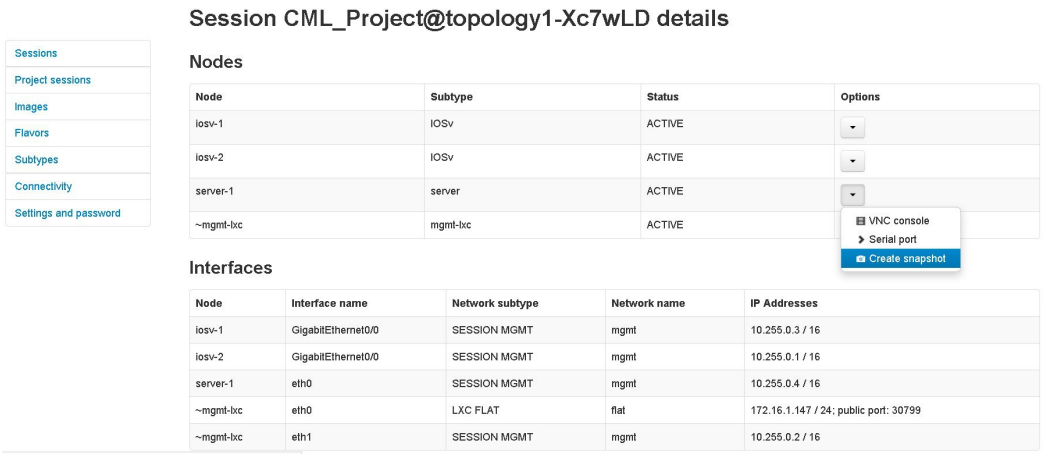


Project details for the newly created snapshot are displayed.

***Figure 61: Newly Created Disk Image***

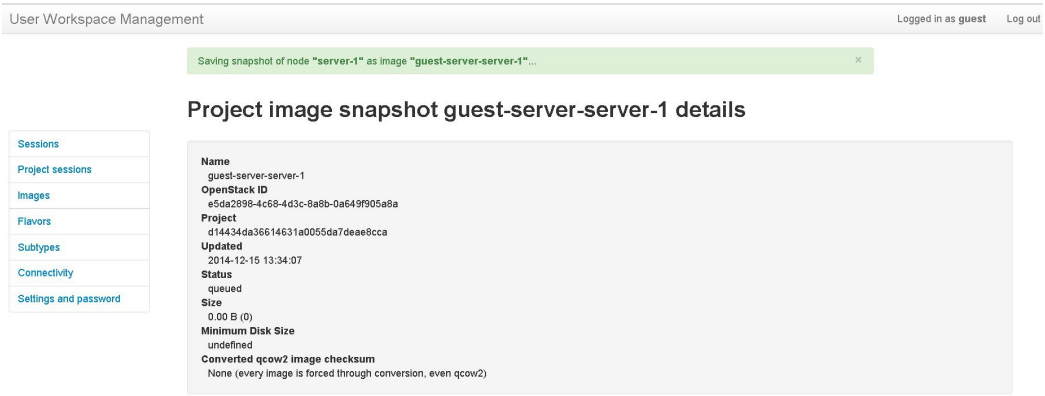

# Reuse the Image Snapshot

To reuse the image snapshot, complete the following steps.

**Step 1** Create a new topology or open an existing topology.

**Step 2** On the canvas, add a node to the topology.

**Step 3** Select the node on the canvas.
The sample topology opens in the **Topology Editor** canvas.

**Step 4** In the **Properties** > **Node** view, click **Browse** beside the **VM Image** field.
The **Select VM Image** dialog box appears.

**Step 5** Select the applicable image snapshot and click **OK**.

***Figure 62: Select the Image Snapshot to Use***

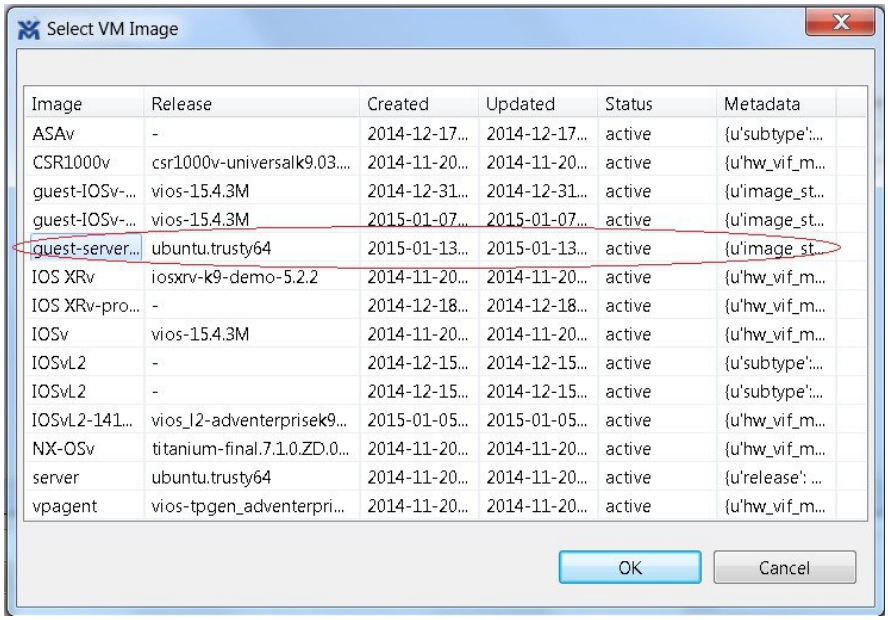

Details for the image snapshot are visible in the **VM Image** field under **Properties** > **Node** view.

***Figure 63: Image Snapshot Selected***

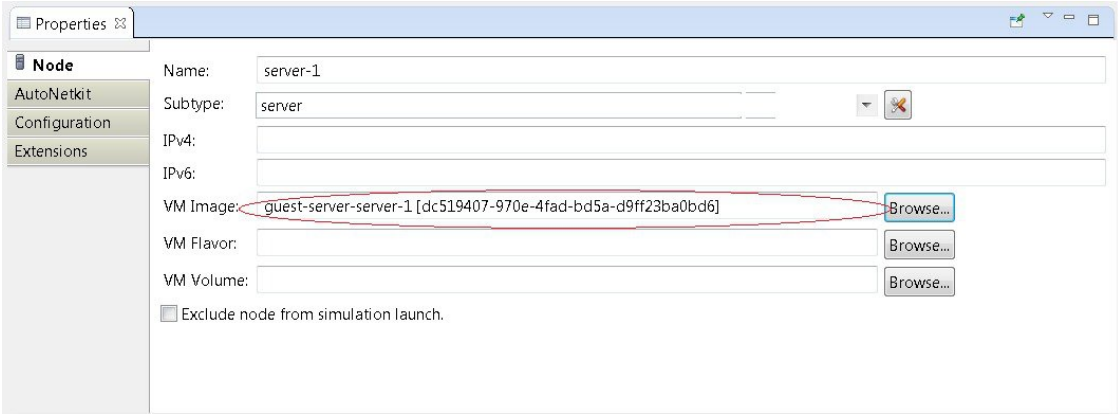


# Latency, Jitter and Packet Loss Control Options

The availability of link-level parameters allows users to investigate and understand the impact on services of transmission characteristics encountered in the physical world. With a running simulation, you are able to select links between the nodes in the simulation and set latency, jitter and packet-loss values on those links. This enables you to create links that have properties seen in the physical world such as transatlantic or transcontinental latencies or packet-loss.

The link parameters can be applied on any link, except for those connected to a FLAT or SNAT external connector. The values set by the user are applied bi-directionally, meaning that setting a latency value of 100ms results in 100ms from node A to node B and 100ms from node B to node A for the return path. That is 200ms in total. The same is true for packet-loss. Ten packets sent from node A on a link with 10% packet-loss results in 9 packets being received on node B. The packet loss will also be applied on the return path meaning that another packet may be lost between node B and node A.

You can set these link-level parameters in one of three ways:

- In the Cisco Modeling Labs Client.

***Figure 64: Setting Link Parameters in the Cisco Modeling Labs Client***

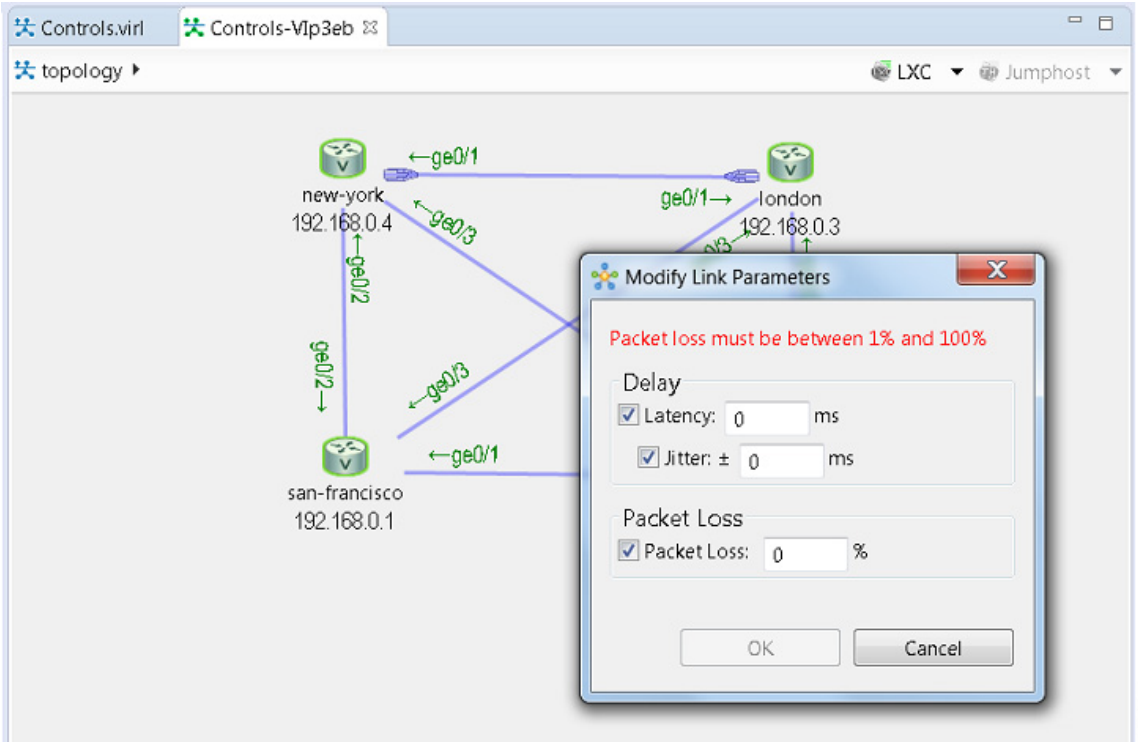

- In the User Workspace Management interface.

***Figure 65: Setting Link Parameters in the User Workspace Management Interface***

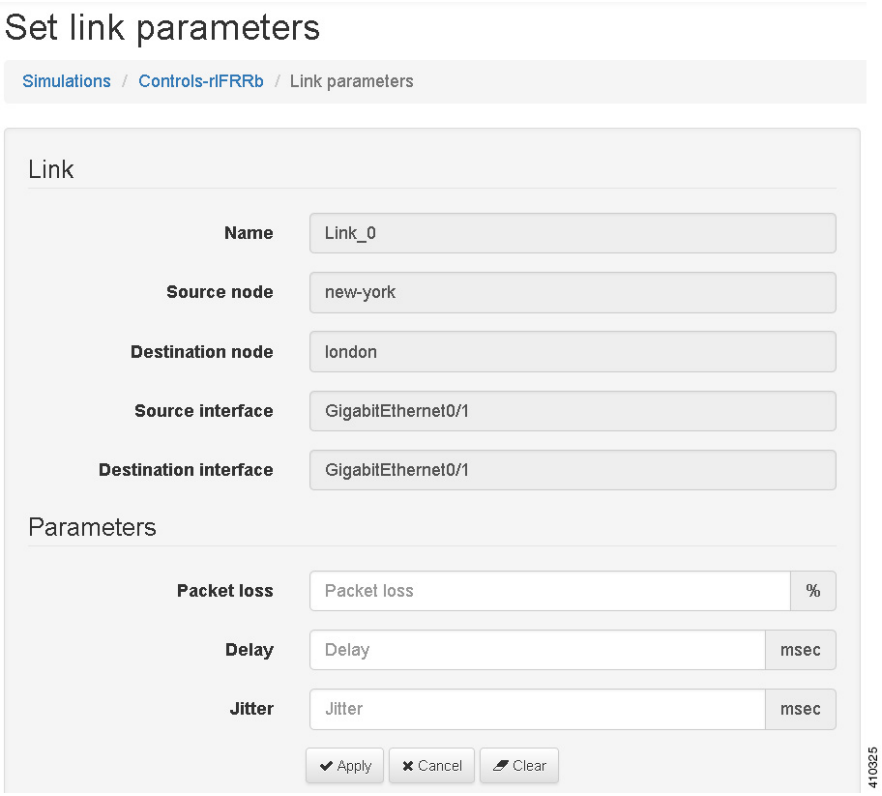

- In the Live Visualization view of the running simulation.
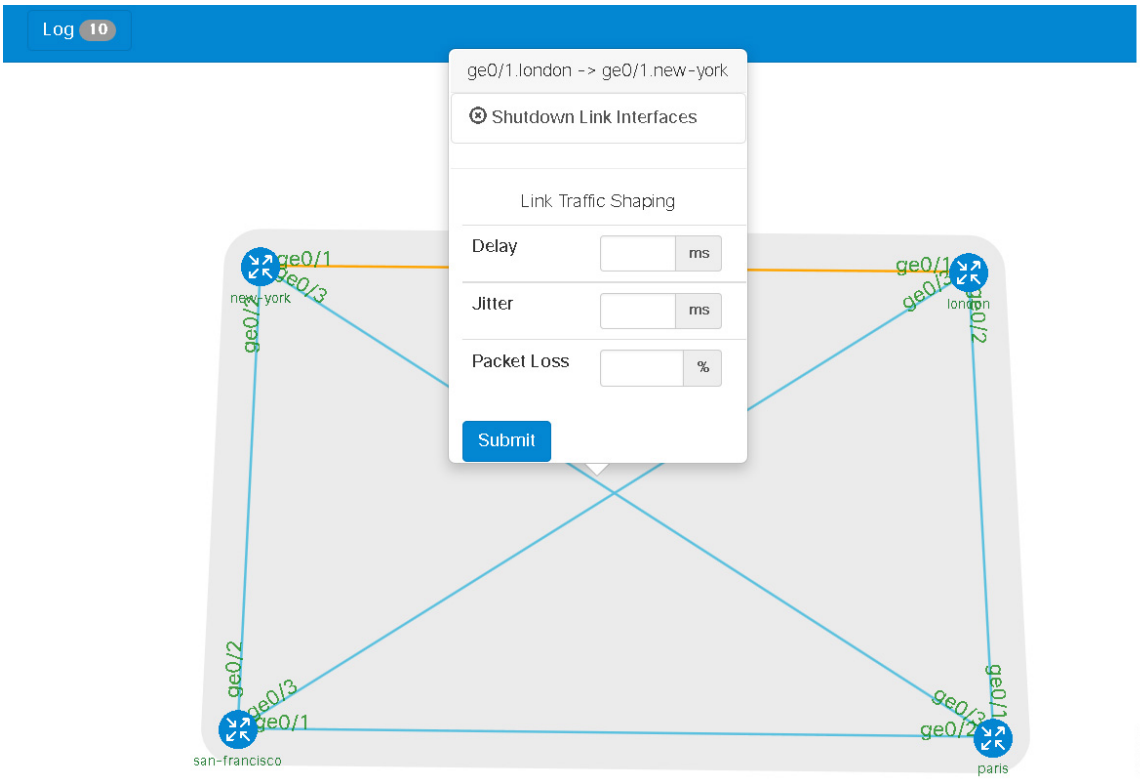
***Figure 66: Setting Link Parameters in the Live Visualization View***